



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 :

H04L 29/06

A1

(11) International Publication Number:

WO 98/56150

(43) International Publication Date:

10 December 1998 (10.12.98)

(21) International Application Number: PCT/US98/11201

(22) International Filing Date: 1 June 1998 (01.06.98)

(30) Priority Data:

60/048,723

2 June 1997 (02.06.97)

US

(71) Applicant: UNISYS CORPORATION [US/US]; Township Line and Union Meeting Roads, P.O. Box 500, Blue Bell, PA 19424-0001 (US).

(72) Inventors: JOHNSON, Robert, A.; 4070 Prospect Hill Lane, Pottstown, PA 19464 (US). EBERSOLE, Dwayne, E.; 179 Hostetter Lane, Lancaster, PA 17602-4139 (US). DISNEY, William, W.; 378 North Buck Road, Downingtown, PA 19335 (US).

(74) Agent: STARR, Mark, T.; Unisys Corporation, Township Line and Union Meeting Roads, P.O. Box 500, Blue Bell, PA 19424-0001 (US).

(81) Designated States: BR, DE, GB, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).


Published

*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(54) Title: SHARED USE OF A NETWORK INTERFACE CARD BETWEEN HETEROGENEOUS COMPUTER SYSTEMS

(57) Abstract

A first network protocol provider (44), executing on a first computer system (100) and having a first network address associated therewith, and a second network protocol provider (58), executing on a second computer system (102) and having a second network address associated therewith, share a same network interface card (50) installed on the second computer system. Apparatus providing this ability comprises an interconnection (62, 64, 66, 70; 65, 67, 70'; 63, 70'') between the first and second computer systems over which data can be transferred between them, and a router (80) executing on the second computer system (102) that routes data among the first network protocol provider (44) (via the interconnection), the second network protocol provider (58), and the network interface card (50) in a manner that enables the network interface card (50) to be shared between the first and second network protocol providers. For example, when data is received from the network (15) by the network interface card (50), the router (80) routes the data to the first network protocol provider (44), via the interconnection, when a network address received with the data matches the first network address, but passes the received data to the second network protocol provider (58) when the network address received with the data matches the second network address. A "loopback" capability (176) also allows data addressed from one network protocol provider to the other to be transmitted between them directly (via the interconnection), bypassing the network interface card (50). The router (80) can also handle broadcast addresses.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

SHARED USE OF A NETWORK INTERFACE CARD BETWEEN HETEROGENEOUS COMPUTER SYSTEMS

5

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the United States Patent & Trademark Office patent file or records, but otherwise reserves all copyright rights
10 whatsoever.

BACKGROUND

Field of the Invention

15 The present invention relates to the field of computer networking, and, more particularly, to apparatus and methods for allowing two computer systems to access a network via a shared network interface card installed on one of them.

Description of the Prior Art

20 The ability for heterogeneous computer systems to communicate with each other over a network using standard and/or proprietary networking protocols is known. Most computer systems have some form of networking architecture that enables the computer system to perform networking in accordance with those protocols. Such a networking architecture typically comprises both system software and hardware. Figure 1 is a block diagram illustrating the
25 components of a networking architecture employed by a Unisys A Series enterprise server 10 in order to communicate with other hosts, or nodes, on a network 15.

The A Series enterprise server 10 executes the Unisys MCP operating system 12, and has an I/O subsystem that comprises one or more I/O Modules (IOM) 14 housed within the A Series chassis. The IOM 14 implements a Unisys proprietary I/O bus architecture referred to as CS-BUS II or CS-Bus III (hereinafter "the CS Bus"). A plurality of card slots, e.g. slots 16a-d, are provided for connecting interface cards, referred to as "channel adapters", into the CS Bus. Different groups, or racks, of channel adapter slots are each controlled by a Channel Manager Unit (CMU) (e.g., CMUs 18a, 18b). An IOM can contain several CMUs, each of which controls a different rack of channel adapter card slots via the CS-Bus. The CMUs manage the physical and data layers of the I/O process.

Channel adapter cards, which each may occupy one or more channel adapter card slots within the IOM 14, provide various connectivity solutions for the A Series enterprise server 10. For example, Unisys provides a channel adapter card that implements the Small Computer System Interface(SCSI) protocol for connecting SCSI peripherals to the enterprise server.

For network connectivity, Unisys provides several channel adapters to support various physical networking protocols. These channel adapters are generally referred to as network processors (NP). For example, Unisys ICP22 and ICP26 network processors are channel adapter cards that implement the Ethernet network protocol and can be used to connect an A Series enterprise server to an Ethernet network. Unisys also provides network processors for connectivity to FDDI and ATM networks. As shown in Figure 1, a number of different network processors (e.g., NPs 20a, 20b, and 20c) can be installed in respective channel adapter slots (e.g., slots 16b, 16c, and 16d) of the IOM 14, in order to provide different network connectivity solutions.

As shown in the more detailed view of network processor 20c (installed in channel adapter slot 16d), a network processor may comprise a plurality of different lines, e.g., Line0, Line1 ... LineN. A line represents a physical endpoint within a network. For example, the Unisys ICP22 network processor has two lines, each of which comprises a separate Ethernet connection - one line could be connected to one Ethernet network, and the other to a different Ethernet network.

Each line of a network processor can have one station group defined on that line. A station group consists of one or more stations. A station is a logical endpoint that represents a logical dialog on that line. Thus, more than one logical dialog can take place over a given line of

a network processor. This is achieved through multiplexing. For example, with a connection-oriented networking protocol, such as the Burroughs Network Architecture - Version 2 protocol (BNAv2), one station may represent a logical dialog with one other BNAv2 host on the network, whereas another station may represent a logical dialog to a different BNAv2 host. As illustrated in Figure 1, for example, Station0 of LineN may represent a logical dialog with BNAv2 host 22, and Station1 of LineN may represent a logical dialog with BNAv2 host 24. For networking protocols that are not connection-oriented, like the Internet Protocol (IP), only one station needs to be defined to handle all communications for that protocol stack. For example, in Figure 1, StationN of LineN could be defined as the logical endpoint for all IP traffic over LineN. A

Local Area Network Station Group (LANSG) module 26, which comprises software executing on the network processor 20c, provides callable procedures for creating and maintaining stations and station groups on the various lines of the network processor 20d and for sending and receiving data over them.

Other software components that execute on the network processor 20d include a Queue Service Provider (QSP) module 28, which handles the multiplexing and demultiplexing of data for all stations defined on a given NP, and two stub modules - a Network Services Manager stub (NSM-stub) 30 and a Link Layer Manager stub (LLM-stub) 32 - which interface with corresponding modules of a Core Network Services (CNS) software component 34, to and from modules within the MCP environment.

Generally, a network processor (e.g., NP 20a, 20b, or 20c) implements the data link and physical layers of the 7-layer ISO Reference Model. Higher level networking protocols that a client application 46 may wish to employ in order to communicate with applications running on different hosts of the network 15, such as the BNAv2 and TCP/IP networking protocols, are implemented as network protocol providers on the A Series system 10. A network protocol provider is a software module that implements these higher level networking protocols. For example, Unisys provides both BNAv2 Host Resident Network Provider (HRNP) modules and TCP/IP HRNP modules. In the example of Figure 1, a BNAv2 HRNP 42 and a TCP/IP HRNP 44 are shown.

The Core Network Services (CNS) software 34 provides support for the network protocol providers 42, 44 and handles the initialization and maintenance of network processors and the station groups defined thereon. Specifically, CNS 34 comprises a Network Services Manager

(NSM) 36 that initializes and manages the network processors (e.g., 20a, 20b, 20c) installed in the system, and a Link Layer Manager (LLM) 38 that initializes and maintains the identity and attributes of each station group defined on a given network processor. Another component (not shown) of CNS 34 validates attributes associated with station groups and stations created on a network processor. These attributes are passed between the network processor and CNS 34 via a control dialog when the stations are defined. Like the stub procedures for the NSM and LLM modules 36, 38, network processors also have a stub procedure (LLAH, not shown) that corresponds to the attribute handler of CNS 34. An NPSUPPORT software library 40, as well as portions of the MCP operating system 12, provide routines and procedure calls that serve as an interface between a network processor and the CNS 34 and network protocol providers 42, 44, and control loading of software to the NPs and dumping of their state.

Each network processor has an associated identifier that uniquely identifies that network processor within the system 10. When a network processor is initialized and brought on-line, the NSM-stub 30 in the network processor interfaces with the NSM 36 of CNS 34 via a control dialog in order to pass its identifier to the NSM 36. The NSM 36 manages the identifiers of all active network processors.

Each station group and station defined for a given network processor also has a unique identifier associated with it. Via a control dialog established between the LLM-stub 32 on the network processor and the LLM 38 of CNS 34, the station and station group identifiers are passed to the LLM 38 during initialization. Within the LLM 38, a station corresponds to a connection, and a station group corresponds to a connection group.

As mentioned above, the ability to define multiple stations (*i.e.*, a station group) on a single physical line of a network processor is achieved through multiplexing. Specifically, the QSP 28 in the network processor multiplexes inbound and outbound data for multiple stations on a given line. Moreover, the QSP is responsible for distributing request and response data between the NSM 36 and NSM-stub 30 and between the LLM 38 and LLM-stub 32. To that end, each entity on the network processor that receives outbound data from the MCP, including every station, the NSM-stub 30, and the LLM-stub 32, is assigned a unique Remote Queue Reference (RQR) by the QSP. The NSM-stub RQR is reported to the NSM 36 within CNS 34 via NPSUPPORT when the NP is loaded. The LLM-stub RQR is reported to the LLM 38 via

the NSM 36 by the NSM-stub 30 when the NP initializes. All of the station RQRs are reported to the HRNPs as the stations open.

When a client application is required to send data via network 15 to some other host or node on the network 15, such as another BNAv2 Host or another TCP/IP host, it invokes the services of the appropriate network protocol provider, *e.g.*, 42, 44. The network protocol provider 42, 44 determines the appropriate network processor and station on which the data is to be output, adds protocol headers, and makes a corresponding request to the MCP 12 that includes the identifier of the network processor and the RQR of the station. The data and associated RQR are passed from the MCP 12 to the QSP 28 on the network processor (*e.g.*, network processor 20c), which, in combination with the LANS module 26, sends the data out to the network 15 via the appropriate line (*e.g.*, Line0, Line1, ... or LineN) as part of the logical dialog represented by the designated station.

When data is received from the network 15 on a given line, the LANS module 26 determines, from header information associated with the data, the station (*i.e.* logical dialog) for which the data is intended. The LANS module 26, QSP modules 28, in combination with portions of the MCP 12 and NPSUPPORT library 40, pass the received data to the appropriate network protocol provider 42, 44 associated with that station, along with an indication of which station received the data. For example, one of the stations on LineN of the network processor 20c of Figure 1 (*e.g.*, station0) may be defined as the logical endpoint for the BNAv2 HRNP 42, while a different station (*e.g.*, station1) may be defined as the logical endpoint on which all IP traffic over LineN is received for the TCP/IP HRNP 44. When a frame of data is received from the network on LineN, the LANS module 26 determines from header information which of the network protocol providers (*i.e.*, stations) is intended to receive the data. This determination is performed in accordance with the methods described in commonly assigned, U.S. Patent No. 5,379,296, entitled "Method and Apparatus for Interfacing a Workstation to a Plurality of Computer Platforms" (Johnson et al.).

In addition to its use in A Series computers, the foregoing networking architecture is also employed in Unisys ClearPath HMP NX enterprise servers. A ClearPath HMP NX server comprises an A Series enterprise server tightly integrated with a server running Microsoft Window NT. Please note that "Microsoft," "Windows," and "Windows NT" are registered trademarks of Microsoft Corporation. Additional information concerning the foregoing

networking architecture can be found in the following documents, each of which is available from Unisys Corporation, assignee of the present invention, and each of which is hereby incorporated by reference in its entirety:

ClearPath HMP NX Series with Windows NT Network Services Implementation Guide
5 (Part No. 4198 6670);

BNA/CNS Network Implementation Guide, Volume 2: Configuration (Part No. 3789
7014);

ClearPath HMP NX Series with Windows NT Implementations and Operations Guide
(Part No. 3807 6542);

10 *ClearPath HMP NX Series with Windows NT Migration Guide* (Part No. 3807 7730);
Networking Capabilities Overview (Part No. 3789 7139)

Networking Operations Reference Manual, Volumes 1 and 2: Commands and Inquiries
(Part No. 3787 7917); and

Networking Products Installation Guide (Part No. 4198 4840).

15 Using a Unisys ICP22 network processor, which is an Ethernet-based channel adapter, it has been possible in the past for a Unisys A Series enterprise server to communicate with a workstation or personal computer (PC) over a network. An example of this ability is illustrated in Figure 2. In this example, the A Series enterprise server 10 communicates with an Intel-based workstation 48 running the Microsoft Windows NT operating system (hereinafter "the NT
20 server"). The A Series enterprise server is connected to the network via network processor 20a, which may, for example, be a Unisys ICP22 Ethernet-based network processor.

The I/O subsystem of the NT server 48 comprises portions of the NT operating system kernel, an EISA or PCI bus 52, and appropriate device driver software. To provide network connectivity, a network interface card (NIC) 50 is installed in an available bus slot on the NT
25 server 48. The NT server may support one or both of the PCI and EISA bus standards. NICs are available for both bus standards.

A NIC device driver 54 that typically is sold with the NIC card 50 is installed in the kernel space of the NT operating system. The NIC device driver 54 interfaces with a higher level network protocol provider, such as an implementation of the TCP/IP protocol. Microsoft
30 Corporation provides an implementation of the TCP/IP protocol in the form of a kernel level device driver, also referred to as a transport protocol driver, named TCPIP.SYS 58. TCPIP.SYS

58 interfaces with the NIC device driver 54 via NDIS, an industry standard Network Driver Interface Specification jointly developed by Microsoft and 3Com. NDIS defines an interface for communication between hardware-independent protocol drivers, such as TCPIP.SYS, which implement the Data Link, Network, and Transport layers of the OSI model, and hardware-
5 dependent NIC drivers which provide an interface to the NIC hardware and which correspond to the Physical Layer of the OSI model. A client program 60 on the NT server can communicate over the network 15 in accordance with the TCP/IP protocol by issuing suitable calls via the NT operating system to the TCPIP.SYS protocol driver.

Network interface cards and associated device drivers for NT servers are available from a
10 number of Original Equipment Manufactures (OEMs). OEM NICs are available at relatively low cost for a variety of different network media standards, including Ethernet, Fast-Ethernet, etc. As new network standards evolve, OEMs are quick to design and produce NICs to support these standards. Because these NICs are developed for industry standard I/O bus architectures, such as EISA and PCI, which are found in the many computer systems today, the economies of
15 scale result in fast cycle development times and extremely low prices for consumers.

On the contrary, it takes significantly longer and costs significantly more to design and produce a new network processor for a proprietary bus architecture, such as the CS-BUS II architecture of Unisys A Series enterprise servers. Vendors of proprietary systems cannot achieve the same economies of scale as the open system NIC vendors, and network processors,
20 or NIC cards, for proprietary systems therefore typically cost significantly more than their open systems counterparts. To avoid the costs associated with the development of NIC cards for proprietary systems, it would be desirable if it were possible for two computer systems to each access a network via a shared network interface card installed on one of the systems. The present invention provides such a capability.

SUMMARY OF THE INVENTION

The present invention is directed to methods and apparatus that enable a first network protocol provider, executing on a first computer system and having a first network address associated therewith, and a second network protocol provider, executing on a second computer
30 system and having a second network address associated therewith, to both send and receive data over a network via a same network interface card installed on the second computer system. In

accordance with a preferred embodiment thereof, the present invention comprises an interconnection that couples the input/output (I/O) subsystem of the first computer system to the I/O subsystem of the second computer system and over which data can be transmitted between systems, and a router executing on the second computer system that routes data among the first network protocol provider (via the interconnection), the second network protocol provider, and the network interface card in a manner that enables the network interface card to be shared between the first and second network protocol providers. Specifically, when data is received from the network by the network interface card, the router routes the data to the first network protocol provider, via the interconnection, when a network address received with the data matches the first network address, but passes the received data to the second network protocol provider when the network address received with the data matches the second network address. Data received from the network with a broadcast address is routed to both the first (via the interconnection) and second network protocol providers. Outgoing data from each network protocol provider is routed to the network via that same network interface card. However, when outgoing data from one of the first and second network protocol providers is addressed to the other network protocol provider, a "loopback" capability enables the router to route that data directly from one network protocol provider to the other, bypassing the network interface card. Also, outgoing data from one of the first and second network protocol providers that has a broadcast address is routed to the network via the network interface card, and also to the other network protocol provider via the loopback path (bypassing the network interface card). In all cases, data routed between the first and second computer systems is transmitted over the interconnection. Additionally, it is understood that the present invention can be used to route data to and from a plurality of network protocol providers on the first computer system and/or a plurality of network protocol providers on the second system, via the same network interface card.

Preferably, the router simulates the operation of a network protocol provider when interfacing with a device driver of the network interface card, and the router simulates the operation of a device driver of the network interface card when interfacing with the first and second network protocol providers, whereby the functionality of the router is transparent to the network interface card and to the first and second network protocol providers.

The interconnection between the I/O subsystem of the first computer system and the I/O subsystem of the second computer system preferably comprises a physical connection between the I/O subsystems over which data can be transmitted between them, and an interconnection device driver on the second computer system that controls access by the second computer system to the physical connection. The interface between the interconnection device driver and other components on the second computer system is preferably implemented in the form of a procedure registration mechanism. In this manner, different interconnection device drivers can be installed on the second computer system for different physical connections, in a manner that is transparent to the other components of the present invention. For example, when the first and second computer systems are separate physical units, the physical connection may comprise suitable hardware (*e.g.*, interface boards) installed in available slots of the I/O buses of each system and a cable that provide a connection between them. Alternatively, where the first computer system is emulated within the second system, the physical connection may be emulated within the second system in the form of a memory-to-memory connection.

One advantage of the present invention is that it can be used to enable a computer system that has a proprietary I/O subsystem to access a network via a standard network interface card developed for a more open, or widely used I/O subsystem such as that found in systems that support the Microsoft Windows NT operating system. This avoids the need to develop network interface cards for the proprietary system as new networking protocols and standards evolve. Instead, the present invention allows the proprietary system to share network interface cards designed for the open system.

Additional features and advantages of the present invention will become evident hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing summary, as well as the following detailed description of the preferred embodiment, is better understood when read in conjunction with the appended drawings. For the purpose of illustrating the invention, there is shown in the drawings an embodiment that is presently preferred, it being understood, however, that the invention is not limited to the specific methods and instrumentalities disclosed. In the drawings:

Figure 1 is a block diagram illustrating the components of a prior art networking architecture employed by Unisys A Series enterprise servers in order to communicate with other hosts, or nodes, on a network;

Figure 2 is a block diagram showing a prior art method by which a Unisys A Series enterprise server can communicate via a network with a server running Microsoft Windows NT;

Figure 3 is a block diagram illustrating one embodiment of apparatus that enables two computer systems to share a single network interface card installed on one of them, in accordance with the present invention;

Figure 4 is a block diagram illustrating an alternative embodiment of an interconnection of the apparatus of Figure 3;

Figure 5 is a block diagram illustrating yet another embodiment of the interconnection of the apparatus of Figure 3;

Figures 6A-F are flow diagrams that further illustrate the operation of the interconnections illustrated in Figures 3-5;

Figure 7 is a diagram illustrating the different data paths over which the router of the present invention may route incoming and outgoing network data;

Figure 8 is a block diagram illustrating further details of an implementation of the router of the present invention, in accordance with a preferred embodiment;

Figures 9A-9C comprise a pseudocode listing illustrating further details of one embodiment of the routing functionality of the router of the present invention, and also one embodiment of a method of the present invention;

Figures 10A-D are flow diagrams illustrating further details of another embodiment of the routing functionality of the router of the present invention, and also another embodiment of the method of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention is directed to methods and apparatus that enable a first network protocol provider, executing on a first computer system and having a first network address associated therewith, and a second network protocol provider, executing on a second computer system and having a second network address associated therewith, to both send and receive data over a network via a same network interface card installed on the second computer system. In

accordance with a preferred embodiment thereof, the present invention comprises an interconnection that couples the input/output (I/O) subsystem of the first computer system to the I/O subsystem of the second computer system and over which data can be transmitted between systems, and a router executing on the second computer system that and a router executing on
5 the second computer system that routes data among the first network protocol provider (via the interconnection), the second network protocol provider, and the network interface card in a manner that enables the network interface card to be shared between the first and second network protocol providers.

In one embodiment, described more fully hereinafter, the methods and apparatus of the
10 present invention may be implemented as part of a Cooperative Networking Platform (CNP) (sometimes also referred to "NX/Network Services" or "NNS") provided as a feature of Unisys ClearPath HMP NX enterprise servers, in which, as mentioned above, a Unisys A Series enterprise server is tightly integrated with an Intel-based server running Microsoft Window NT. In that embodiment, the A Series enterprise server comprises the first computer system and the
15 NT server comprises the second computer system. As embodied in that environment, the present invention allows a network protocol provider (*e.g.*, BNAv2 HRNP or TCP/IP HRNP) on the A Series server to send and receive data over a network via a standard network interface card installed in the NT server. A network protocol provider (*e.g.*, TCPIP.SYS) on the NT server can also access the network via that NIC, and thus, the NIC is shared by both systems. This
20 capability is advantageous because NICs for NT servers are widely available at relatively low cost. The need to develop proprietary NICs (*i.e.*, network processors) for the A Series server, and the significant costs associated therewith, can be avoided.

The methods and apparatus of the present invention, or certain aspects or portions thereof, may take the form of program code (*i.e.*, instructions) embodied in tangible media, such
25 as floppy diskettes, CD-ROMs, hard drives, or any other machine-readable storage medium, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. The methods and apparatus of the present invention may also be embodied in the form of program code that is transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via
30 any other form of transmission, wherein, when the program code is received and loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the

invention. When implemented on a general-purpose processor, the program code combines with the processor to provide a unique apparatus that operates analogously to specific logic circuits.

Referring now to the drawings, wherein like numerals represent like elements throughout, Figure 3 is a block diagram illustrating one embodiment of the present invention, in which the methods and apparatus of the present invention are implemented as part of a Cooperative Networking Platform (CNP) deployed on a Unisys ClearPath HMP NX computer system ("the ClearPath system"). As shown, the ClearPath system comprises a Unisys A Series enterprise server 100 and an Intel-based server 102 running Windows NT 102 ("the NT server"). In this embodiment, the A Series enterprise server 100 defines a first computer system and the NT server 102 defines a second computer system. A first network protocol provider 44 is provided on the A Series system 100, in this case a TCP/IP HRNP, and it has a network address (*i.e.*, IP address) associated with it. A second network protocol provider 58 is provided on the NT server 102, in this case TCPIP.SYS (available from Microsoft Corporation), and it has its own network address (*i.e.*, IP address) associated with it that defines a second network address in this embodiment. Other network protocol providers may be installed on the A Series and NT servers as well. For example, on the A Series server, a BNAv2 HRNP 42 may be provided. Note, however, that because the BNAv2 protocol is a Unisys proprietary protocol, the BNAv2 HRNP 42 does not have an IP address associated with it. A network interface card (NIC) 50 is installed in a slot of the I/O bus (*e.g.*, EISA or PCI) of the NT server 102. In this embodiment, any LAN-type NIC that is compatible with Windows NT can be employed. Preferably, the NIC supports the Fast-Ethernet networking protocol (*e.g.*, 100Base-T). NICs of this type are available from numerous vendors and original equipment manufacturers (OEMs). NICs supporting other physical media types, such as Ethernet/802.3, FDDI, or Gigabit Ethernet, can alternatively be employed. Typically, a NIC vendor will supply a device driver with the NIC, which is installed in the kernel space of the operating system so that other entities on the system can access the networking functionality of the NIC. The NIC 50 of the exemplary system of Figure 3 has a device driver 54 ("`<nicdrv>.sys`") that is installed in the Windows NT kernel space, as shown.

According to the present invention, apparatus and methods are provided that enable both the first and second network protocol providers 44, 58 to access the network 15 via the same

NIC 50 installed on the second computer system 102. That is, the present invention enables the NIC 50 to be shared between the two systems 100, 102.

Apparatus of the present invention comprises an interconnection that couples the I/O subsystem of the A Series server 100 to the I/O subsystem of the NT server 102 so that data can be transmitted between the two servers, and a router 80 that executes on the second computer system (*i.e.*, the NT server 102). The router 80 routes data among the first network protocol provider (via the interconnection), the second network protocol provider, and the NIC 50 so that the NIC 50 can be shared by both network protocol providers. More specifically, when data is received from the network 15 by the NIC 50, the router 80 routes the data to the first network protocol provider (*e.g.*, TCP/IP HRNP 44), via the interconnection, when a network address received with the data matches the first network address, but passes the received data to the second network protocol provider (*e.g.*, TCPIP.SYS 58) when the network address received with the data matches the second network address. Data received from the network 15 with a broadcast address is routed to both the first (via the interconnection) and second network protocol providers. Outgoing data from each network protocol provider is routed to the network 15 via that same NIC 50. However, when outgoing data from one of the first and second network protocol providers is addressed to the other network protocol provider, a "loopback" capability enables the router 80 to route that data directly from one network protocol provider to the other, bypassing the NIC 50. Also, outgoing data from one of the first and second network protocol providers that has a broadcast address is routed to the network 15 via the shared NIC 50, and also to the other network protocol provider via the loopback path (bypassing the NIC 50).

The present invention is not limited to use with only one network protocol provider on each system. Rather, the present invention can be used to route data to and from multiple network protocol providers on each system, each network protocol provider having its own network address. Incoming data is routed to the appropriate network protocol provider based on the received network address. Moreover, in the present embodiment, in addition to being able to route data to TCP/IP network protocol providers (*e.g.*, TCP/IP HRNP 44) on the A Series server 100, the router 80 is also capable of routing incoming data addressed to network protocol providers that support the BNAv2 networking protocol (*e.g.*, BNAv2 HRNP 42), as further described hereinafter.

Additional details of the interconnection and of the router 80 are provided hereinafter.

I. THE INTERCONNECTION

As mentioned above, the interconnection of the apparatus of the present invention couples the I/O subsystem of the A Series server 100 to the I/O subsystem of the NT server 102 to provide a relatively high speed data path between systems. Preferably, the interconnection comprises a physical connection between the I/O subsystems of the first and second computers, and an interconnection device driver 70 that controls access to the physical connection by other software modules on the NT server.

A. One Embodiment

In the embodiment of Figure 3, the physical connection comprises a feedthrough card 62 installed in a channel adapter slot of the A Series server 100, an EISA Personal Computer Channel Adapter (EPCCA) card 66 installed in an EISA slot of the I/O bus of the NT server 102, and a CS-BUS II cable that connects the CS-BUS II of the A Series system to the EPCCA card 66 via the feedthrough card 62. The interconnection device driver (ICD) 70 is installed in the kernel space of the NT operating system. It controls access to the physical connection (specifically the EPCCA card 66) by other modules on the NT server.

Although not necessary in every embodiment of the present invention, in the embodiment illustrated in Figure 3, the interconnection further comprises several modules that are analogous to similarly named components in the traditional Unisys networking architecture illustrated in Figure 1 and described in the background section of this specification. These modules include a Queue Service Provider module 76 that functions analogously to the QSP 28 of Figure 1, a LANSg module 78 that functions analogously to the LANSg module 26 of Figure 1, and NSM-stub and LLM stub modules 84, 86 that function analogously to the corresponding components 30, 32 of Figure 1. In addition, LDM and LLAH modules 82, 88 are provided which function analogously to the similarly named components (not shown) in a traditional Unisys networking architecture. In combination with the physical connection (*i.e.*, feedthrough card 62, cable 64, and EPCCA card 66) and the interconnection device driver 70, these modules simulate a traditional channel adapter-based network processor of the type described above and illustrated in Figure 1. In this manner, the features and advantages of the present invention are achieved with relatively little modification of the networking components (*e.g.*, CNS 34, NPSUPPORT 40, MCP 12, etc.) on the A Series server 100. With the exception of the LDM 82,

NSM-stub 84, LLM-stub 86, and LLAH 88 modules, the main components of the interconnection are implemented as Windows NT kernel-level device drivers to avoid unnecessary data copies that would otherwise occur when transferring data from the NT kernel space to the user space. Each of the foregoing components and modules of the interconnection is described below in more detail.

1. EPCCA card 66

The EISA Personal Computer Channel Adapter (EPCCA) card 66 plugs into an available EISA bus slot in the NT server 102. The EPCCA card 66 bridges the EISA bus on the NT server 102 to the CS-BUS II of the A Series server 100 via the cable 64 and feedthrough card 62. The feedthrough card 62 plugs into an available channel adapter slot in the IOM 14, providing a direct connection to the data and control lines of the CS-BUS II. Microcode on the EPCCA card 66 emulates an A Series channel adapter.

2. Interconnection Device Driver 70 (PCCA 72 and OPENCA 74)

In the present embodiment, the interconnection device driver 70 comprises a PCCA device driver 72 and an OPENCA device driver 74. The PCCA driver 72 initializes, terminates, configures, and communicates with the EPCCA hardware/microcode 66. The PCCA driver 72 interacts with the OPENCA driver 74 via a procedural interface. This procedural interface passes 32-byte I/O Messages (IOMs) between the EPCCA board 66 and OPENCA 74.

The OPENCA driver 74 acts as an intermediary between PCCA 72 and the rest of the components of the interconnection, as well as providing controller functions for the datapath. It interprets commands received from the A Series server 100 via the PCCA driver 72 and EPCCA card 66, generates results, and sets up memory descriptors for moving data between the NT server 102 and the A Series server 100. OPENCA interfaces with the QSP module 76 via a procedural interface; the IOM-based API used by the PCCA driver 66 is hidden from the QSP 76 and other modules of the interconnection.

3. Queue Service Provider (QSP) 76

The QSP 76 provides a message queuing function that is necessary to interface the NSM and LSM-stubs 84, 86, and the LANGS module 78 to their peers on the A Series server 100 via NPSUPPORT 40. The QSP 76 functions analogously to the QSP 28 in a traditional network processor, as described above and illustrated in Figure 1. Specifically, the QSP 76 and NPSUPPORT 40 exchange control messages to initialize, establish and terminate pipes.

configure, and report errors. They also build and parse headers placed at the beginning of data messages. These headers specify message blocking, message length, and remote queue references (RQRs). As mentioned above, RQRs are used to multiplex many point-to-point dialogs over the lines managed by the LANSNG module. Entities on the A Series server 100 that correspond to the stations managed by the LANSNG module are assigned unique RQRs, as are the stations to which they correspond. These RQRs are then exchanged out of band via NSM/NSM-Stub and LLM/LLM-Stub control dialogs. Once the RQRs have been exchanged, the originator of a message inserts the other side's RQR in the header, so the receiver can assign that message to the appropriate queue. Thus, the QSP 76 enables multiple dialogs to be multiplexed over the physical connection between the A Series and NT servers.

4. LAN Station Group (LANSNG) 78

Like the QSP 76, the LANSNG 78 functions analogously to the LANSNG 26 in a traditional network processor, as described above and illustrated in Figure 1. LANSNG 78 controls the initialization and termination of station groups on particular lines, as well as the individual stations of each group. It also controls the setting and retrieval of attributes for those groups, the reporting of asynchronous indications, and the passing of data messages thereon. In this embodiment, the LANSNG 78 communicates via a STREAMS interface to the QSP 76.

STREAMS is an industry standard device driver interface popularized on UNIX systems, but also available for Windows NT systems. In the present embodiment, the STREAMS interface is implemented on the NT server 102 using the Mentat Portable Streams for Windows NT (MPS) product available from Mentat, Inc., 1145 Gayley Ave. Suite 315, Los Angeles, CA 90024 USA.

Data from the NT server (e.g., data received from the network 15 via a shared NIC 50) that is destined for a network protocol provider on the A Series server 100 is forwarded to the LANSNG 78 and then sent through the QSP 76, the interconnection device driver (ICD) 70, and the physical connection to NPSUPPORT 40 on the A Series server 100. In the present embodiment, the LANSNG module 78 manages physical lines in a manner similar to the LANSNG module 26 of a traditional Unisys network processor. In this embodiment, however, the lines that the LANSNG module 78 manages are implemented by network interface cards installed in the NT server 102. For example, the shared NIC 50 installed in the NT server 102 defines one line managed by the LANSNG module 78. More than one shared NIC 50 can be installed in the NT server 102, each defining a different line of the Unisys networking architecture. Moreover, as

described in co-pending application serial no. _____, entitled "A Virtual Lan Interface For High-Speed Communications Between Heterogeneous Computer Systems." (Attorney Docket: USYS-0038/TN086), a simulated FDDI local area network within the memory space of the NT server defines another line within the LANSg module 78. As described in the co-

5 pending application, the simulated FDDI network, in combination with the interconnection described herein, provides another connectivity solution between the TCP/IP network protocol providers 44, 58 on each server. In the present embodiment, the simulated FDDI local area network is always defined as Line0.

The LANSg module 78 maintains a mapping of line numbers to AdapterNames. Line

10 numbers are assigned to station groups by configuration commands entered on the A Series server 100. The LLM communicates these line numbers to the LLM-stub 86. AdapterNames are assigned by Windows NT as NICs are configured into the NT server 102 and stored in the Windows NT Registry. LANSg 78 will obtain the AdapterNames of those NICs that are bound to LANSg 78 in the Registry and will maintain a mapping to the associated line number and

15 other information. The following table illustrates the information maintained for each line number in use:

Line Number	Adapter Name	Adapter Handle	Media
0	always simulated FDDI LAN	obtained from call to <i>NdisOpenAdapter</i>	always FDDI
1	obtained from the registry		Ethernet
2	-	-	-
3	-	-	-
4	-	-	-
5	-	-	-
6	-	-	-
7	-	-	-
8	-	-	-
9	-	-	-
10	-	-	-
11	-	-	-
12	-	-	-
13	-	-	-
14	-	-	-
15	-	-	-

5. Load/Dump Module (LDM) 82

The LDM 82 provides a mechanism for simulating a load of a network processor firmware file (a process that is performed during initialization of traditional network processors) to the CNP platform on the NT server 102 and for initiating a CNP state dump. The LDM resides inside an executable file, CNP.EXE, that is initiated as an NT service when the NT server 102 boots up. As part of the simulated firmware loading, the LDM 82 initiates the NSM-stub process 84, which in turn initiates the LLM-stub process 86, and initiates the QSP 76.

6. Network Services Module stub (NSM-stub) 84

The NSM-stub module 84 is also part of the CNP.EXE executable and is responsible for carrying out commands/responses sent to it by the NSM 36 (a component of CNS 34) on the A Series server 100. Essentially, it performs the functions of the NSM-stub 30 of a typical Unisys network processor. In the present embodiment, the NSM-Stub 84 interfaces with the QSP module 76 via a STREAMS interface using standard STREAMS calls (*i.e.*, open, close, ioctl, putmsg, getmsg).

7. Link Layer Manager stub (LLM-stub) 86

The LLM-stub module 86, also part of CNP.EXE, is responsible for carrying out commands/responses sent to it by the LLM 38 (a component of CNS 34) on the A Series server 100. Essentially, it performs the functions of the LLM-stub 32 of a typical Unisys network processor. In the present embodiment, the LLM-Stub 86 likewise interfaces with the QSP module 76 via a STREAMS interface using standard STREAMS calls (*i.e.*, open, close, ioctl, putmsg, getmsg).

8. Link Layer Attribute Handler (LLAH) 88

The LLAH module 88, another module within CNP.EXE, functions similarly to its counterpart in a traditional Unisys network processor. Specifically, the LLAH module 88 is responsible for performing the detailed processing associated with parsing, validating, and building attribute lists. The LLAH enforces attribute range validity rules and checks for inter-attribute consistency. The LLM-Stub 86 is the only module that interfaces with the LLAH.

B. Host Interface Function (HIF) - Alternative Embodiments

Still referring to Figure 3, the interconnection device driver 70, including its PCCA and OPENCA drivers 72, 74 in the present embodiment, and the physical connection formed by the feedthrough card 62, cable 64, and EPCCA card 66, together define a Host Interface Function

(HIF). According to another feature of the present invention, the procedural interface between the QSP 76 and the interconnection device driver 70 of the HIF is designed to isolate the QSP 76 from the HIF. This enables the present invention to be employed with different implementations of the HIF. Specifically, the procedural interface between the QSP 76 and the interconnection device driver 70 is established through a process by which each module publishes entry points (i.e., pointers) to the procedures that implement its functionality, along with any required variable values. Another device driver entity called NNSDRLOG.SYS (not shown) maintains a record of these entry points.

The interconnection device driver 70 of the HIF registers the following entry points and attributes:

HifSendBlockToHost() -- a function called by QSP to deliver a block of data to the MCP;

HifOpenUnit() -- a function called by QSP to initialize one of several pipes (units) through which data received from the LANS module 78 can be transmitted to the appropriate entity on the A Series server 100;.

HifCloseUnit() -- a function called by QSP to indicate that one of the pipes (units) is terminated;

maxQueuesSupported -- a variable initialized by the HIF which QSP can reference to determine how many pipes (queues/units) it can use to send messages to the MCP 12 of the A Series server 100; and

platform -- a variable initialized by the HIF which identifies (via an enumeration) a particular implementation of the HIF (two alternate HIF implementations are described below and illustrated in Figures 4 and 5, respectively). In the present embodiment, these functions and variables are implemented by the OPENCA driver 74 of the interconnection device driver 70.

Likewise, the QSP 78 registers the following entry points:

QspAckBlockToHost() -- a function called by the ICD to indicate to the QSP 78 that a particular block of messages has been successfully delivered to the MCP 12;

QspReset() -- a function called by the ICD to indicate to the QSP 78 that communications with the MCP 12 via the interconnection have been lost, and that pending messages should be flushed; and

QspLRPut() -- a function called by the ICD to deliver a block of data from the A Series server to the QSP 78.

In order to invoke one of these functions, a call is made to the registered entry point for that function. As a result of this indirection, different interconnection device drivers can be
5 installed for different implementations of the HIF in a manner that is completely transparent to the QSP 76.

Figures 4 and 5 illustrate two alternate embodiments of the HIF, which illustrate the modularity provided by the procedural interface design described above. In Figure 4, the physical connection (*i.e.*, the feedthrough card 62, cable 64, and EPCCA card 66) is replaced by
10 a PCI Bridge card 67 that connects via a cable 65 directly to a port on one of the CMU's 18b of the IOM 14 of the A Series server 100. By connecting directly to the CMU, some of the latency inherent in the CS-Bus II protocol is avoided. This provides a more direct, higher speed connection between the I/O subsystems of the two servers 100, 102. Because the physical connection is changed, a modified interconnection device driver 70' is provided. The modified
15 interconnection device driver 70' comprises a single device driver module, PXN 73, that provides the interface between the QSP 76 and the hardware on the PCI Bridge card 67. However, the procedural interface, and the mechanism by which the QSP 76 and interconnection device driver 70' register entry points to the respective procedures of that interface is unchanged. Accordingly, the changes to the HIF are transparent to the QSP 76 and the other modules of the
20 present invention that comprise the Cooperative Networking Platform (CNP).

Figure 5 is an embodiment in which the A Series system is emulated through software in the NT server 102. Unisys provides such an emulated system in its ClearPath HMP NX 4200 series enterprise servers. In this embodiment, the physical connection is emulated such that it becomes a memory-to-memory connection 63 between the memory space of the emulated I/O
25 subsystem 14' and the memory space of the NT system. The emulated connection 63 functions in a manner similar to the feedthrough card 62, cable 64, EPCCA card 66, and PCCA 72 components of the hardware implementation of Figure 3. The interconnection device driver 70' in this embodiment comprises a modified form 74' of the OPENCA module 74 of the implementation of Figure 3. Again, however, the procedural interface between the modified
30 OPENCA module 74' and the QSP 76 is not changed, so that the emulated A Series server and

its emulated connection 63 to the NT server is transparent to the QSP 76 and the other modules of the present invention that comprise the Cooperative Networking Platform (CNP).

C. Operation

Figures 6A-6F provide further details of how data is transferred between the A Series server 100 and the NT server 102 via the interconnection device driver of the HIF and the QSP module 76. The details provided in Figures 6A-6E are applicable to any of the three embodiments of the HIF shown in Figures 3, 4, and 5. Thus, as used in the following discussion, the term interconnection device driver (ICD) refers to any of the three interconnection device driver embodiments described above.

The QSP 76 multiplexes multiple client dialogs (*e.g.*, dialogs with the NSM-stub and LLM-stub modules 84, 86 and with the different stations defined by LANSG 78) over one or more transfer units. Units are an abstraction of the communication paths supported by the interconnection device driver (ICD). Units may be logical dialogs or physical devices. In order to more fully utilize the unit resources, the QSP 76 may aggregate messages waiting for transfer over a same unit into a block that can be transferred in a single operation. The QSP 76 supports such blocking by providing a Message-Count field in its message headers. The first message header in a block contains the number of messages that the block contains in its Message-Count field. Subsequent message headers within the block have a zero value in that field.

The ICD then takes each block and programs the physical connection (*i.e.*, the EPPCA board 66, the PCI Bridge card 67, or the emulated memory-to-memory connection 63, depending upon the implementation) to transfer the block to the A Series server. In the reverse direction, the ICD is awakened when a message is transferred via the physical connection into the memory of the NT server 102, either by an interrupt (in the case of the hardware connections of Figures 3 and 4) or by a function call (in the case of the emulated connection 63 of Figure 5). The ICD delivers the received message to the QSP 76, which in turn, distributes it to the appropriate client dialog (*e.g.*, NSM-stub 84, LLM-stub 86, or a given station defined by LANSG 78), based on the RQR associated with the message.

Figures 6A-6D provide further information concerning the steps performed by the QSP 76 and ICD in transferring messages from a client on the NT server 102 (*e.g.*, NSM-stub 84, LLM-stub 86, or a station defined by LANSG 78) to the A Series server 100 via the physical connection. This transfer process begins when a client, for example, the LANSG module 78,

which may need to pass data received from the NIC 50 to the A Series server 100, calls the QSP 76 requesting that a message (e.g., the data received from the network) be transferred to the A Series server 100. A parameter is passed with the request that points to non-contiguous message segments that comprise the full message. At step 112, the QSP 76 determines on what unit the message should be transferred. Next, at step 114, the QSP 76 calculates the total size of the message by examining each non-contiguous segment in the message. At step 116, a header is added to the beginning of the message, and a descriptor list is built that points to the header and to each segment in the message. Next, at step 118, the QSP 76 determines whether blocking (described above) is supported for this unit. If so, at step 120, the QSP 76 determines whether any blocks are presently waiting for transfer. If so, at step 121, the QSP 76 determines whether the message will fit in the last pending block. If so, then at step 122, the QSP 76 adds the descriptor list to the last pending block. Control then passes to step 127 (Fig. 6B).

If in step 118, blocking is not supported for this unit, or if in step 120 it is determined that there are no blocks presently waiting for transfer, or if in step 121 it is determined that the message will not fit in the last pending block, then control passes in all three cases to step 124. At step 124, the QSP 76 builds a block containing only the descriptor list built in step 116. Next, at step 126, the newly created block is added to the list of pending blocks. Control then passes to step 127 (Fig. 6B).

Referring now to Figure 6B, at step 127, the QSP 76 determines whether any blocks are pending. If not, the QSP 76 simply returns to the client. If however, there are pending blocks to be transferred, then control passes to step 128.

At step 128, the QSP 76 attempts to send the first block in the list of pending blocks to the ICD by invoking the `HifSendBlockToHost()` procedure of the ICD. As indicated by the arrow labeled "A", the ICD begins processing the request at that point. The steps performed by the ICD are illustrated in Figure 6C. Still referring to Figure 6B, however, the QSP's processing continues to step 130, where the QSP 76 determines whether the ICD accepted the block for transfer. If so, that block is removed from the pending list at step 132, and control loops back to step 127 where the QSP 76 again checks whether there are any pending blocks to be transferred and processing continues for any such subsequent blocks. If, however, in step 130 it is determined that the ICD did not accept a given block for transfer, then the QSP 76 returns to the client, leaving the block containing the message to be sent on the pending list.

Referring now to Figure 6C, the ICD begins processing the `HifSendBlockToHost()` request from the QSP at step 134, where it determines whether the physical connection is in flow-control mode. Flow-control mode is a mode in which the MCP operating system of the A Series server 100 is not prepared to receive data on the specific unit, for example, because no
5 buffer is available. If the physical connection is in flow-control mode, the ICD returns a value of "FALSE" to the QSP 76 and stops processing the transfer at this point. If the physical connection is not in flow-control mode, then control passes to step 136 where the ICD determines whether the physical connection supports a Gather function. Gather is the ability to transfer data from non-contiguous memory regions in one operation. If the physical connection
10 does not support a Gather capability, control passes to step 138 where the ICD copies the data pointed to by the descriptor list (passed to it by the QSP 76) into a contiguous buffer. Next, at step 140, the ICD builds a pseudo descriptor list that points at the single, contiguous buffer. Control then passes to step 142.

At step 142, whether entered directly from step 136 (Gather supported) or from step 140
15 (Gather not supported), the ICD programs the physical connection (*i.e.*, the EPCCA board 66, the PCI Bridge card 67, or the emulated memory-to-memory connection 63 depending upon the particular embodiment) to transfer the data pointed to either by the descriptor list received from the QSP 76 (Gather) or the pseudo descriptor list created in step 140 (no Gather). The ICD then returns a value of "TRUE" to the QSP 76.

Figure 6D illustrates the steps performed by the ICD and QSP 76 when the transfer
20 completes. As shown, when the transfer completes, the ICD is awakened. At step 144, the ICD receives an indication of whether the transfer completed successfully. If not, control passes to step 146 where the ICD attempts to recover from the error by, for example, retransferring the block in question, resetting the physical connection, etc. If the transfer completed successfully,
25 control passes to step 148. At step 148, the ICD adjusts the flow-control state of the physical connection. This is done because in the embodiments of the physical connection described above, the interconnection is polled. When a transfer completes, the interconnection may not be able to initiate another transfer until it is polled again, so the flow-control state is adjusted to reflect this. Next at step 150, the ICD calls the `QspAckBlockToHost()` procedure to notify the
30 QSP that the transfer is complete and to indicate which descriptor list was transferred. At step 152, the ICD performs a cleanup procedure and then returns.

As shown at point "B", when the QSP 76 receives the QspAckBlockToHost() indication from the ICD, notifying it that the transfer completed successfully, the QSP 76 enters step 154 where all messages in the transferred block are released, causing the clients that sent them to be notified that they were successfully transferred. At step 156, the block structures, including the message headers and descriptor list, are recycled and made available for subsequent transfers. Control then loops back to step 127 of Figure 6B for processing of subsequent blocks.

Figures 6E-6F illustrate the steps performed by the ICD and QSP 76 in transmitting a message from the A Series server 100 to the NT server 102. As shown, prior to the reception of any messages from the A Series server via the physical connection, the ICD makes empty receive buffers available to the connection. When a message is transferred from the A Series server 100 to the NT server via the physical connection (e.g., through the feedthrough card 62, across cable 64, and through the EPCCA card 66 in the embodiment of Figure 3), the ICD is awakened with an indication that a message has been received into one of the empty receive buffers that it posted. At step 158, the ICD passes the message to the QSP 76 using the QspLRPut() function and returns.

At step 160, the QSP 76 determines whether the message is a control message. If so, at step 164, the QSP 76 processes the control message locally, and then releases the message at step 166 and returns. If the message is not a control message, then control passes to step 162. At step 162, the QSP 76 determines from the RQR in the message header which station is to receive the message. Next, at step 168, the message is passed to the appropriate station.

Referring now to Figure 6F, when the QSP 76 or one of its clients releases the message buffer, a free message callback function of the ICD is invoked. At step 170, the ICD adds the freed buffer to the list of available buffers which the physical connection can then use to receive subsequent messages in the manner described above.

II. THE ROUTER 80

Referring again to Figure 3 for purposes of illustration, as mentioned above, the router 80 routes data among a first network protocol provider (e.g., BNAv2 HRNP 42 or TCP/IP HRNP 44) on a first computer system (via the interconnection), a second network protocol provider (e.g., TCPIP.SYS 58) on a second computer system, and a NIC 50 on the second computer system, so that the NIC 50 can be shared by both network protocol providers.

Figure 7 illustrates the various data paths that the router 80 provides in accordance with a preferred embodiment of the present invention. When data is received from the network 15 by the NIC 50, the router 80 routes the data to the first network protocol provider (e.g., TCP/IP HRNP 44 or BNAv2 HRNP 42), via the interconnection, when a network address received with the data matches the first network address, as illustrated by path 174. When the received data has an address that matches that of the second network protocol provider (e.g., TCP/IP.SYS 58), the data is routed to that network protocol provider, as illustrated by path 172. Data received from the network 15 with a broadcast address is routed to both the first (via the interconnection) and second network protocol providers, via paths 172 and 174. Outgoing data from each network protocol provider is routed to the NIC 50 via the respective paths 172, 174. However, when outgoing data from one of the first and second network protocol providers is addressed to the other network protocol provider, a "loopback" capability enables the router 80 to route that data directly from one network protocol provider to the other, bypassing the NIC 50, as illustrated by path 176. Also, outgoing data from one of the first and second network protocol providers that has a broadcast address is routed to the shared NIC 50 via the respective path 172 or 174, and also to the other network protocol provider via the loopback path 176, bypassing the NIC 50. In the present embodiment, the "loopback" path 176 is only implemented for TCP/IP-to-TCP/IP transfers.

As also illustrated in Figures 3 and 7, the present invention is not limited to use with only one network protocol provider on each system. Rather, the present invention can be used to route data to and from multiple network protocol providers on each system, each network protocol provider having its own network address. Moreover, in the present embodiment, network protocol providers on the A Series server 100 may include both TCP/IP network protocol providers (e.g., TCP/IP HRNP 44) and BNAv2 network protocol providers (e.g., BNAv2 HRNP 42). In this embodiment, the router 80 distinguishes between BNAv2 frames and TCP/IP frames according to the method described in commonly assigned, U.S. Patent No. 5,379,296, entitled "Method and Apparatus for Interfacing a Workstation to a Plurality of Computer Platforms" (Johnson et al.), which is hereby incorporated by reference in its entirety.

Figure 8 is a block diagram illustrating further details of the implementation of the router 80 in accordance with the preferred embodiment. In this embodiment, the LANS 78, QSP 76 are both implemented as kernel-level device drivers within the kernel space of the NT operating

system. The interface between LANS 78 and QSP 76 is a STREAMS interface implemented using the Mentat Portable Streams (MPS) for Windows NT run-time environment. In this embodiment, the functionality of the router 80 is implemented in two components, a transport driver component (JTD) 182 and a NIC driver component (JND) 184. The JND component 184 is implemented in the form of an NDIS Miniport NIC driver, referred to herein as JANUS .SYS. "Janus" is a reference from Roman mythology to the ancient god of gates and doorways, depicted with two faces looking in opposite directions. The JTD component is implemented as a module (JTD.C) that is separately compiled from the LANS.C module (which contains the LANS functionality), but linked as part of LANS.SYS.

The JTD component 182 of the router 80 maintains an adapter table with the state of each NIC for which it is responsible, including the IP addresses of the network protocol providers on the A Series and NT servers 100, 102 to which it may have to route incoming data in accordance with the present invention. The following is a C-language structure definition of the adapter table:

```
typedef struct {
    NDIS_STRING  adapter_name;
    NDIS_HANDLE  jtd_adapter_handle; // returned from NdisOpenAdapter
    NDIS_HANDLE  jnd_adapter_handle; // JND's Adapter Handle
    UINT         medium;             // returned from NdisOpenAdapter
    NDIS_HANDLE  packet_pool_handle; // returned from NdisAllocatePacketPool
    NDIS_HANDLE  buffer_pool_handle; // returned from NdisAllocateBufferPool
    U32          mcp_ip_addr;        // A Series IP address for this adapter
    U32          nt_ip_addr;         // NT server IP address for this adapter
    sg_tab_t     *sg_tab_p;          // ptr to stn group assoc. w/ this
                                     adapter
} jtd_adapter_info_t;
jtd_adapter_info_t adapter_table[MAX_ADAPTERS_SUPPORTED];
```

The "adapter_name" field of the structure contains the AdapterName assigned to the particular shared NIC (e.g., NIC 50) that the entry represents. The "jtd_adapter_handle" and "medium" fields contain values returned by the NdisOpenAdapter function of the NDIS specification. The "packet_pool_handle" and "buffer_pool_handle" fields contain values returned by the NdisAllocatePacketPool and NdisAllocateBufferPool functions of the NDIS specification. The "jnd_adapter_handle" field contains a value that is obtained from the JND component 184. This

value represents the unique handle used by JND to identify the adapter. The "mcp_ip_addr" field contains the IP address of the network protocol provider (e.g., TCP/IP HRNP 44) on the A Series server that shares this particular NIC, and the "nt_ip_addr" field contains the IP address of the network protocol provider (e.g., TCPIP.SYS 58) on the NT server 102 that shares this NIC.

- 5 Alternatively, other embodiments may support multiple IP addresses on a single adapter line, for a given network protocol provider.

Even though JTD.C is linked with LANS.G.C into LANS.G.SYS and could store this info in the station group table of LANS.G, a separate (but linked) adapter table is chosen in this design to isolate the NDIS interface in JTD.C and to minimize the size of the station group table (a "mask" of station group table size is used in the interface from the LLM-Stub and LANS.G.C for setting station group attributes). The adapter table is linked to the station group table of LANS.G 78 by the following additional entry in the station group table which points to a given entry of the adapter table:

```
15 typedef struct {
    . . .
    jtd_adapter_info_t *adapter_info_p;
    . . .
    sg_table_t;
20 ..

```

- The JTD component 182 obtains the IP addresses of the TCP/IP network protocol providers on each server 100, 102 to which it may have to route incoming data as follows. On the NT server 102, when TCPIP.SYS 58 is configured to use an IP address on a particular NIC (e.g., NIC 50), it stores an entry containing that IP address in the Windows NT Registry at the following location:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<AdapterName>\
Parameters\TCPIP.
```

- The JTD component 182 has access to this location in the Registry for all of the NICs that are configured to be "shared" and will store the IP addresses retrieved from the Registry for each shared NIC in its local tables.

On the A Series server 102, IP addresses are configured using the NW TCPIP TCPIPID command. This command causes the TCP/IP HRNP on the A Series server 100 to associate a single IP address with a network processor and a line on that network processor. As mentioned above, in the present embodiment, the interconnection simulates a traditional network processor

from the viewpoint of the A Series server. Line0 is reserved for the simulated FDDI network described in the aforementioned co-pending application serial no. _____, entitled "A Virtual Lan Interface For High-Speed Communications Between Heterogeneous Computer Systems" (Attorney Docket: USYS-0038/TN086). Line numbers above Line0 can be used for

5 NICs that are to be shared in accordance with the present invention. LANSg maps line numbers to particular adapters as described previously. When the TCPIP HRNP 44 is configured with a particular IP address, it sends the configured IP address to the LLM-stub 86 using a Set Station Interface (SSI) command, along with other miscellaneous data. LANSg 78 receives the SSI command from the LLM-Stub 86, and passes the IP address to the JTD component 182 of the

10 router 80 for inclusion in its tables.

LANSg 78 also maintains a mapping of the IP addresses associated with a particular shared NIC in its tables. The following table illustrates the information maintained for this mapping:

Adapter Handle	MCP IP Address	NT IP Address

15 As shown, for each shared NIC (identified by its Adapter Handle), a mapping is maintained to both the IP address of the network protocol provider on the A Series server 100 that shares that particular NIC, and the IP address of the network protocol provider on the NT server that shares that NIC.

The NIC 50 and its device driver (<nicdrv>.sys) 54 are installed in the normal fashion in

20 accordance with the instructions supplied by the NIC vendor. However, the bindings in Windows NT for the NIC 50 (as well as for any other shared NIC) must be changed from their original settings when the NIC and its driver are installed. Specifically, the Janus NIC driver (JND) and the Janus transport driver (JTD) components of the router 80 need to be inserted into the bindings. All NT transport drivers (e.g., TCPIP.SYS) must be re-bound to "adapters" in the

25 JND. In addition, the JTD must bind as a transport driver to the NIC driver 54. Since JTD is implemented as a module that is linked into LANSg.SYS, LANSg.SYS is actually installed as the transport driver.

In addition to the foregoing, the following new entry is required in the Windows NT Registry - EnableIPRoutingInMCP. This Registry entry is used by the JTD component to designate the TCP/IP HRNP 44 on the A Series server 100 as an IP router for frames with "unknown" IP address and to designate the relationship of the IP router on the A Series server 100 to the IP router on the NT server 102 (if enabled). EnableIPRoutingInMCP may have 3 possible values: "Primary", "Secondary", and "Disabled". If set to "Primary", the JTD always directs any inbound IP or ARP frames with unknown IP addresses to the TCP/IP HRNP 44 on the A Series server 100, via the interconnection. If EnableIPRoutingInMCP is set to "Secondary", then the JTD will only direct any inbound IP or ARP frames with unknown IP address to the TCP/IP HRNP 44 on the A Series server 100, if EnableIPRouting (an existing NT Registry entry) is not set. If EnableIPRouting is set, then it takes precedence and IP or ARP frames with unknown IP addresses will be sent to TCPIP.SYS 58 on the NT server 102. If EnableIPRoutingInMCP is set to "Disabled", then JTD will not direct any inbound IP or ARP frames with unknown IP address to the TCP/IP HRNP 44, but rather will send such frames to TCPIP.SYS 58.

On the A Series server, changes have been made to the CNS 34 software. First, two new attributes have been defined for LAN-type connection groups, as follows:

1. ADAPTERNAME - This is a settable/gettable attribute that is of type "string", with a maximum length of 17 characters. This attribute is used to map a given connection group defined within CNS 34 to the particular shared NIC (e.g., NIC 50) that it represents.

2. ADAPTERTYPE - This is a settable/gettable attribute that identifies the type of adapter to which the ADAPTERNAME attribute refers. In the present embodiment, the ADAPTERTYPE attribute can be set to VLAN (referring to the simulated FDDI network described in the aforementioned co-pending application) or to FASTETHERNET (referring to the types of NICs supported in this embodiment).

On a shared NIC, the router 80 is transparent to both NT TCPIP.SYS 58 and the TCP/IP HRNP 44 on the A Series server 100. NT TCPIP.SYS interfaces with the router 80 as it would normally with a NIC driver through the NDIS interface. The TCP/IP HRNP 44 on the A Series server 100 also interfaces with the router 80 as it would normally with a network interface via LANS 78. Neither NT TCPIP.SYS 58 nor the A Series TCP/IP HRNP 44 is aware that they

are sharing the same NIC. In addition, the NIC driver 54 interfaces to the router 80 as it would normally with an NDIS protocol driver, such as NT TCPIP.SYS.

Figures 9A-9C provide further details of the inbound routing functionality of the router 80, in accordance with one embodiment thereof, and also illustrate one embodiment of a method of the present invention. The embodiment illustrated in Figures 9A-9C does not, however, support the "loopback" path described above. The functionality is presented in pseudocode form.

Referring to Figure 9A, in line 200, the router 80 determines from the length field of the MAC header of a frame of data received from the network whether length of the frame is less than or equal to 1500 bytes, in which case the frame is assumed to be an 802.3 frame containing an 802.3 MAC header, an 802.2 LLC header, and an 802.1a SNAP (SubNetwork Attachment Point) header. If that is the case, then at line 202, the router 80 determines from the DSAP field of the LLC header of the frame whether the frame is an SNAP frame. If so, then at line 204, the router 80 determines from the control field of the LLC header whether the frame is also an Unnumbered Information (UI) type frame (*i.e.*, a frame type used for "information" (data) frames in the 802.2 LLC protocol. If the frame is not a UI frame, then it is identified as having an UNKNOWN type, and a corresponding value is stored in a `type_of_frame_rcvd` variable at line 208 to indicate the UNKNOWN type. If in line 204, the frame is determined to be a UI frame, then at lines 206, the router 80 further determines whether the frame is an IP frame, an ARP frame, or an UNKNOWN type, and stores the identified type in the `type_of_frame_rcvd` variable.

If in line 202, the frame is not determined to be a SNAP frame, then at lines 210, the router 80 determines from the DSAP field of the LLC header and from a `lookup_source_addr()` function in LANSg 78 (which searches the list of LANSg stations looking for a match between a remote address in the station table entry and the source MAC address in the frame), whether the frame is intended for the BNAv2 HRNP 42 on the A Series server 100. If so, then at line 212, the frame is identified as such by storing a corresponding value in the `type_of_frame_rcvd` variable. If the frame is not identified as a BNAv2 frame, then the frame is assigned an UNKNOWN type at one of the lines 214.

If in step 200, the incoming frame does not have a length less than or equal to 1500 bytes, then it is assumed to be either an Ethernet frame or a large BNAv2 frame. At line 216, the

router 80 checks the DSAP field of the LLC header of the frame to determine if it could be a BNAv2 frame. If so, then at line 218, the router 80 invokes a lookup_source_addr() function, passing it the source_addr field of the MAC header of the frame, which determines from that information whether the frame is addressed to the BNAv2 HRNP 42 on the A Series server 100.

- 5 If so, the lookup_source_addr() function returns a pointer to the station that represents a connection to that BNAv2 HRNP. Then, in line 220, the type_of_frame_rcvd variable is set to "BNA." If, however, no BNAv2 HRNP is identified in line 218, then the frame is assigned the UNKNOWN type in line 222.

- 10 If in line 216, the frame is determined not to have the characteristics of a BNAv2 frame, then it should be either an IP frame or an ARP frame. In lines 224, the router 80 determines from the type field of the Ethernet MAC header which of these frame types has in fact been received. If the type cannot be determined, the frame is simply identified as an UNKNOWN type.

- 15 Once the type of the received frame has been determined, the frame must be routed to the appropriate network protocol provider. In the case of an IP or ARP frame, the dest_addr field of the IP header of the frame is first examined to determine if the frame is to be "broadcast" to all IP addresses on the network (line 226). If so, then an ip_addr_owner variable is set to indicate that both the network protocol provider on the A Series server 100 (e.g., TCP/IP HRNP 44) and the network protocol provider on the NT server 102 (e.g., TCPIP.SYS 58) are to receive the
- 20 frame. If the frame is not a broadcast frame, it is first assumed that it is addressed to neither network protocol provider 44, 58 (line 228). In lines 230, the router 80 determines which of the network protocol providers 44, 58 is to receive the frame, and sets the ip_addr_owner variable accordingly. In lines 232, the frame is then routed to the appropriate network protocol provider based on the value of ip_addr_owner. For example, if the frame is to be routed to the TCP/IP
- 25 HRNP 44 on the A Series server, the router 80 passes the frame to the LANSG 78 module, which, in combination with the other components of the interconnection, passes the frame to the A Series server 100 via the physical connection for forwarding to that network protocol provider. If, instead, the frame is to be received by the network protocol provider (TCPIP.SYS 58) on the NT server 102, then the router 80 passes the frame to that network protocol provider
- 30 via appropriate calls to the NDIS interface. If ip_addr_owner indicates that neither network protocol provider was determined to be the addressee of the frame, then the router 80 routes the

frame based on the values of the EnableIPRoutingInMCP and EnableIPRouting entries in the Windows NT Registry (described above).

In the case of a BNAv2 frame (lines 234), the router 80 passes the frame to the LANS module 78 along with the pointer to the station that represents a connection to the BNAv2
5 HRNP that is to receive the frame. In the case of an UNKNOWN frame type (lines 236), the frame is passed to the NDIS interface for handling of unknown frame types.

Figures 10A-D comprise flow diagrams that illustrate another, preferred embodiment of the routing functionality of the router 80, as well as a preferred embodiment of the method of the present invention. In this embodiment, the "loopback" feature is supported. Figures 10A-10B
10 illustrates the steps performed by the router 80 when a frame of data is received from the network 15 via the shared NIC 50. Figure 10C illustrates the steps performed by the router 80 when a frame of data is sent by a network protocol provider on the A Series server 100 for output to the network 15 via the shared NIC 50. Figure 10D illustrates the steps taken when a frame of data is sent by a network protocol provider on the NT server 102 for output to the
15 network 15 via the shared NIC 50.

Referring to Figures 10A-10B, at step 240, the NIC 50 receives an incoming frame from the network 15. At step 242, the NIC driver 54 calls the NDIS interface to notify it that a frame has been received. At step 244, NDIS calls all Protocol Drivers that are "bound" to the shared NIC 50. In the present embodiment, the JTD component 182 of the router 80 is such a bound
20 protocol driver. As explained above, the JTD component 182 appears to NDIS as a typical protocol driver.

At step 246, the frame is copied into a buffer owned by the LANS module 78. Next, at step 252, the router 80 checks the Destination MAC Address in the frame to determine whether it matches the broadcast address. If so, control passes to step 258 (Fig. 10B). If the Destination
25 MAC Address is not the broadcast address (step 252), then at step 254, the router 80 determines whether the Destination MAC address matches the MAC Address of the NIC 50. If not, then at step 255, the router 80 checks whether the Destination MAC Address is in the Multicast Address List of the NIC. If not, then the frame is discarded at step 256. If the Destination MAC Address is in the Multicast Address List of the NIC, then control passes to step 258 (Fig. 10B).

30 If in step 254, the Destination MAC Address is found to match the MAC Address of the NIC, then control passes to step 248, where the router 80 (specifically the JTD component 182

in the embodiment illustrated in Figure 8) determines from the frame whether it is intended for the BNAv2 HRNP 42 on the A Series system. If so, the frame is passed by LANSg 78 to that BNAv2 HRNP via the HIF at step 250. If the frame is not intended for the BNAv2 HRNP 42, then control passes to step 258 of Figure 10B.

5 Referring now to Figure 10B, at step 258, the router 80 determines whether the frame is an ARP frame. If so, then it is either an ARP request or an ARP reply (determined at steps 260 and 262, respectively). If the frame is an ARP Request, then a copy of the frame is routed to the TCPIP.SYS 58 driver on the NT server 102 at step 268, and the frame itself is routed to the TCP/IP HRNP 44 on the A Series system via the HIF at step 270. If the ARP frame is an ARP
10 Reply, then at step 264, the router 80 determines whether the target protocol address matches the IP address of the TCP/IP HRNP 44 on the A Series server 102. If so, the frame is routed to that network protocol provider via the HIF at step 270. If not, the frame is routed via NDIS to the TCPIP.SYS driver 58 on the NT server 102 at step 266.

If the frame is not an ARP frame, but is an IP frame (step 272), then steps 274 through
15 278 are performed to determine whether the Destination IP Address of the IP frame is either the broadcast address (step 274), the network IP address (step 276), or a Class D Multicast address (step 278). If the destination IP address of the IP frame is any of these, then a copy of the frame is routed to both the TCPIP.SYS driver 58 on the NT server 102 via NDIS (step 268) and to the TCP/IP HRNP 44 on the A Series server 100 via the HIF (step 270). If in step 272, the frame is
20 determined not to be an IP frame, then control passes to step 266 where the frame is routed only to the NT TCPIP.SYS 58.

If the destination IP address is neither of the types checked in steps 274 through 278, then control passes to step 280 where the router 80 determines whether the destination IP address matches that of the TCP/IP HRNP 44 on the A Series system. If so, then the frame is
25 routed to that network protocol provider via the HIF at step 270. If not, then at step 282, the router 80 checks the EnableIPRoutingInMCP entry in the Windows NT Registry. If that entry is set to "Primary", then the frame is routed to the TCP/IP HRNP 44 on the A Series system via the HIF at step 270. If the entry is set to "Disabled", then the frame is routed to the TCPIP.SYS driver on the NT server 102 via NDIS at step 266. If the EnableIPRoutingInMCP entry is set to
30 "secondary", then control passes to step 284, where the router 80 checks the EnableIPRouting entry in the Registry. If the EnableIPRouting entry is set to "True," then the frame is routed to

the TCPIP.SYS driver 58 on the NT server 102 via NDIS at step 266. If the EnableIPRouting entry is set to "False," then the frame is routed to the TCP/IP HRNP 44 on the A Series server 100 via the HIF at step 270.

Figure 10C illustrates the steps performed by the router 80 when a block of data containing frames to be output over the NIC 50 are sent via the HIF from a network protocol provider (e.g., either a BNAv2 HRNP or a TCP/IP HRNP) on the A Series server 100. At step 290, a block of data from the network protocol provider is received on the NT server 102 via the HIF. At step 292, the block of data is queued to the LANS module 78. As indicated at step 294, steps 296 through 308 are applied to each frame in the received block.

For a given frame, at step 296, the LANS module 78 determines whether the frame is from a BNAv2 HRNP or a TCP/IP HRNP. If the frame is a BNAv2 frame, then at step 298, the LANS module 78 prepends BNA headers to the frame in accordance with the BNAv2 protocol, prepends the link level headers, and then delivers the frame to NDIS for output over the NIC 50 (step 308).

If instead, the frame is from a TCP/IP HRNP, then at step 299, the LANS module 78 prepends TCP/IP and link level headers to the frame. Next, at step 300, the router 80 determines whether the destination MAC address is the broadcast address. If so, then at step 302 the router 80 routes a copy of the frame to the TCPIP.SYS driver 58 on the NT server 102 via NDIS, as if the frame had been received by the NIC 50. Next, at step 308, the frame is also routed to the network 15 via the shared NIC 50.

If the destination MAC address of the frame is not the broadcast address, but instead matches the source MAC address, then at step 306, the router 80 routes the frame to the TCPIP.SYS driver 58 on the NT server 102 via NDIS, as if the frame had been received by the NIC 50. If the destination MAC address of the frame matches neither the broadcast address nor the source MAC address, then the frame is simply routed to the network 15 via the shared NIC 50 at step 308. Note that, in the above description, steps 302 and 306 represent the "loopback" capability of the router 80 of the present invention.

Figure 10D illustrates the steps performed when a frame is sent for output by the network protocol provider (e.g., TCPIP.SYS 58) on the NT server 102. The outgoing frame is passed via a call to the NDIS interface at step 320. At step 322, NDIS delivers the frame to the router 80 (specifically the JND component 184 in the embodiment of Figure 8). At step 324, the JND

component 184 passes the frame to the JTD component 182, which, in the present embodiment, is part of the LANS module 78. Next, at step 326, the frame is copied into a buffer provided by LANS 78.

Next, at step 327, the link level headers are stripped from the frame. At step 328, the
5 router 80 determines whether a connection to the TCP/IP HRNP on the A Series server 100, via the HIF, is open. If not, then the router 80 (specifically the JTD component 182) simply routes the frame to the NIC 50 via NDIS at step 338.

If at step 328, a connection to the TCP/IP HRNP on the A Series server 100 is identified, then at step 330, the router 80 determines whether the destination MAC address of the frame
10 matches the broadcast address. If so, then at step 332, a copy of the frame is routed to the TCP/IP HRNP on the A Series server 100 via the HIF (as if it had been received by the NIC 50), and the frame itself is routed to the NIC 50 via NDIS at step 338. If the destination MAC address of the frame does not match the broadcast address, then at step 334, the router 80 determines whether the destination MAC address matches the source MAC address. If so, then
15 at step 336, the frame is only routed to the TCP/IP HRNP on the A Series server 100 via the HIF. If the destination MAC address of the frame matches neither the broadcast address nor the source MAC address, then the frame is only routed to the NIC 50 via NDIS at step 338. Again, note that, in the above description, steps 332 and 336 represent the "loopback" capability of the router 80 of the present invention.

20 As the foregoing illustrates, the present invention is directed to methods and apparatus that enable a first network protocol provider, executing on a first computer system and having a first network address associated therewith, and a second network protocol provider, executing on a second computer system and having a second network address associated therewith, to both send and receive data over a network via a same network interface card installed on the second
25 computer system. It is understood that changes may be made to the embodiments described above without departing from the broad inventive concepts thereof. For example, while the present invention is described above in the context of a system comprising an A Series server and an NT server, it is understood that the methods and apparatus of the present invention can be employed with any two computer systems, whether of the same or different types. Also, while
30 in the most preferred embodiment, the router 80 is described as comprising a JTD component and a JND component, other embodiments may implement the router functionality as a different

combination of components or as a single component. Additionally, the interconnection of the present invention is not limited to the particular embodiments disclosed. Rather, the term "interconnection" is intended to encompass other methods and apparatus for transferring data between the I/O subsystems of the first and second computer systems. For example, other
5 embodiments may not require the functionality of the QSP and LANSg components. Rather, a more direct interface between the interconnection device driver (ICD) and the router could be employed. Accordingly, the present invention is not limited to the particular embodiments disclosed, but is intended to cover all modifications that are within the spirit and scope of the invention as defined by the appended claims.

WHAT IS CLAIMED IS:

1. Apparatus enabling a first network protocol provider, executing on a first computer system and having a first network address associated therewith, and a second network protocol provider, executing on a second computer system and having a second network address associated therewith, to both access a network via a same network interface card installed on the second system, comprising:
 - an interconnection between an input/output (I/O) subsystem of the first computer system and an I/O subsystem of the second computer system, over which data can be transmitted between systems; and
 - a router executing on the second computer system that passes data received by said network interface card to the first network protocol provider, via the interconnection, when a network address received with the data matches the first network address, and that passes the received data to the second network protocol provider when the network address received with the data matches the second network address.
2. The apparatus recited in claim 1, wherein the router further operates to direct data from both the first network protocol provider and the second network protocol provider to the network via said same network interface card in a manner transparent to both network protocol providers.
3. The apparatus recited in claim 1, wherein when one of said first and second network protocol providers addresses data to the other, the router further operates to pass that data to said other network protocol provider directly, bypassing the network interface card in a manner transparent to both network protocol providers.
4. The apparatus recited in claim 1, wherein when data is received via the network interface card that is intended to be broadcast to all network protocol providers, the router detects such data and passes it to both the first and second network protocol providers.

5 5. The apparatus recited in claim 1, wherein when one of said first and second network protocol providers sends data to be output that has a broadcast address, the router operates to route that data to the network via said same network interface card and also to route that data to the other network protocol provider via the interconnection.

5

 6. The apparatus recited in claim 1, wherein the router simulates the operation of a network protocol provider when interfacing with a device driver of the network interface card, and the router simulates the operation of a device driver of the network interface card when interfacing with the first and second network protocol providers, whereby the functionality of the router is transparent to the network interface card and to the first and second network protocol providers.

10

 7. The apparatus recited in claim 1, wherein the second system comprises an NT server and wherein the network interface card has a device driver installed on the second system that controls access to the NIC , and further wherein the NIC device driver and the network protocol provider on the NT server interface to the router via the Network Driver Interface Specification (NDIS).

15

 8. The apparatus recited in claim 1, wherein the router creates and maintains a table containing the network addresses of the first and second network protocol providers.

20

 9. The apparatus recited in claim 1, wherein the first network protocol provider comprises an implementation of the TCP/IP protocol on the first computer system, and wherein the second network protocol provider comprises an implementation of the TCP/IP protocol on the second computer system.

25

 10. The apparatus recited in claim 1, wherein the first network protocol provider comprises a plurality of network protocol providers each having a network address associated therewith.

30

11. The apparatus recited in claim 10, wherein the first network protocol provider comprises a TCP/IP network protocol provider and a BNAv2 network protocol provider.

5 12. The apparatus recited in claim 1, wherein the interconnection between the first computer system and the second computer system comprises:
a physical connection between the I/O subsystem of the first computer system and the I/O subsystem of the second computer system over which data can be transmitted between them;
and
an interconnection device driver on the second computer system that controls access by
10 other modules on the second computer system to the physical connection.

13. The apparatus recited in claim 12, wherein the interconnection between the first and second computer systems further comprises a queue service provider on the second computer system that multiplexes multiple dialogs between systems over the physical
15 connection.

14. The apparatus recited in claim 12 wherein the first computer system comprises a Unisys A Series server and the second computer system comprises an NT server, and wherein the physical connection comprises a connection between a CS-Bus interface of the I/O
20 subsystem of the A Series server and an I/O bus of the I/O subsystem of the NT server.

15. The apparatus recited in claim 14 wherein the I/O bus of the I/O subsystem of the NT server comprises an EISA bus.

25 16. The apparatus recited in claim 12 wherein the first computer system comprises a Unisys A Series server and the second computer system comprises an NT server, and wherein the physical connection comprises a connection between a port of a Channel Manager Unit (CMU) of the I/O subsystem of the A Series server and an I/O bus of the I/O subsystem of the NT server.

17. The apparatus recited in claim 16, wherein the I/O bus of the I/O subsystem of the NT server comprises a PCI bus.

18. The apparatus recited in claim 12 wherein the first computer system, including a
5 memory space of the first computer system, and its I/O subsystem are emulated within the second computer system, and wherein the physical connection is emulated as a memory-to-memory connection between the memory space of the emulated first computer system and the memory space of the second computer system.

10 19. The apparatus recited in claim 12 wherein the interconnection device driver interfaces with another module on the second computer system via a procedural interface in which procedures of the interconnection device driver and procedures of the other module are registered so as to isolate the functionality of the interconnection device driver from the other module in order to permit different physical connections to be employed with correspondingly
15 different interconnection device drivers in a manner that is transparent to said other module.

20. In a system comprising a first computer system and a second computer system, wherein a first network protocol provider executes on the first computer system and has a first network address associated therewith, and a second network protocol provider executes on the
20 second computer system and has a second network address associated therewith, and further wherein an interconnection between the first and second computer systems enables data to be transferred between them, a method for enabling both said first and second network protocol providers to access a network via a same network interface card installed on the second system, said method comprising:

25 routing data received by said network interface card to the first network protocol provider, via the interconnection, when a network address received with the data matches the first network address; and

routing the received data to the second network protocol provider when the network address received with the data matches the second network address.

21. The method recited in claim 20, further comprising routing data from both the first network protocol provider and the second network protocol provider to the network via said same network interface card in a manner transparent to both network protocol providers.

5 22. The method recited in claim 20, wherein, when one of said first and second network protocol providers addresses data to the other, the method further comprises routing that data to said other network protocol provider directly, via the interconnection, thereby bypassing the network interface card in a manner transparent to both network protocol providers.

10 23. The method recited in claim 20, wherein, when data is received via said same network interface card that is intended to be broadcast to all network protocol providers on the network, the method further comprises routing such data to both the first network protocol provider, via the interconnection, and the second network protocol provider.

15 24. The method recited in claim 20, wherein, when one of said first and second network protocol providers sends data to be output that has a broadcast address, the method further comprises routing that data to the network via said same network interface card and also routing that data to the other network protocol provider via the interconnection.




Figure 1
(Prior Art)




Figure 2
(Prior Art)




Figure 3

4/20




Figure 4

5/20


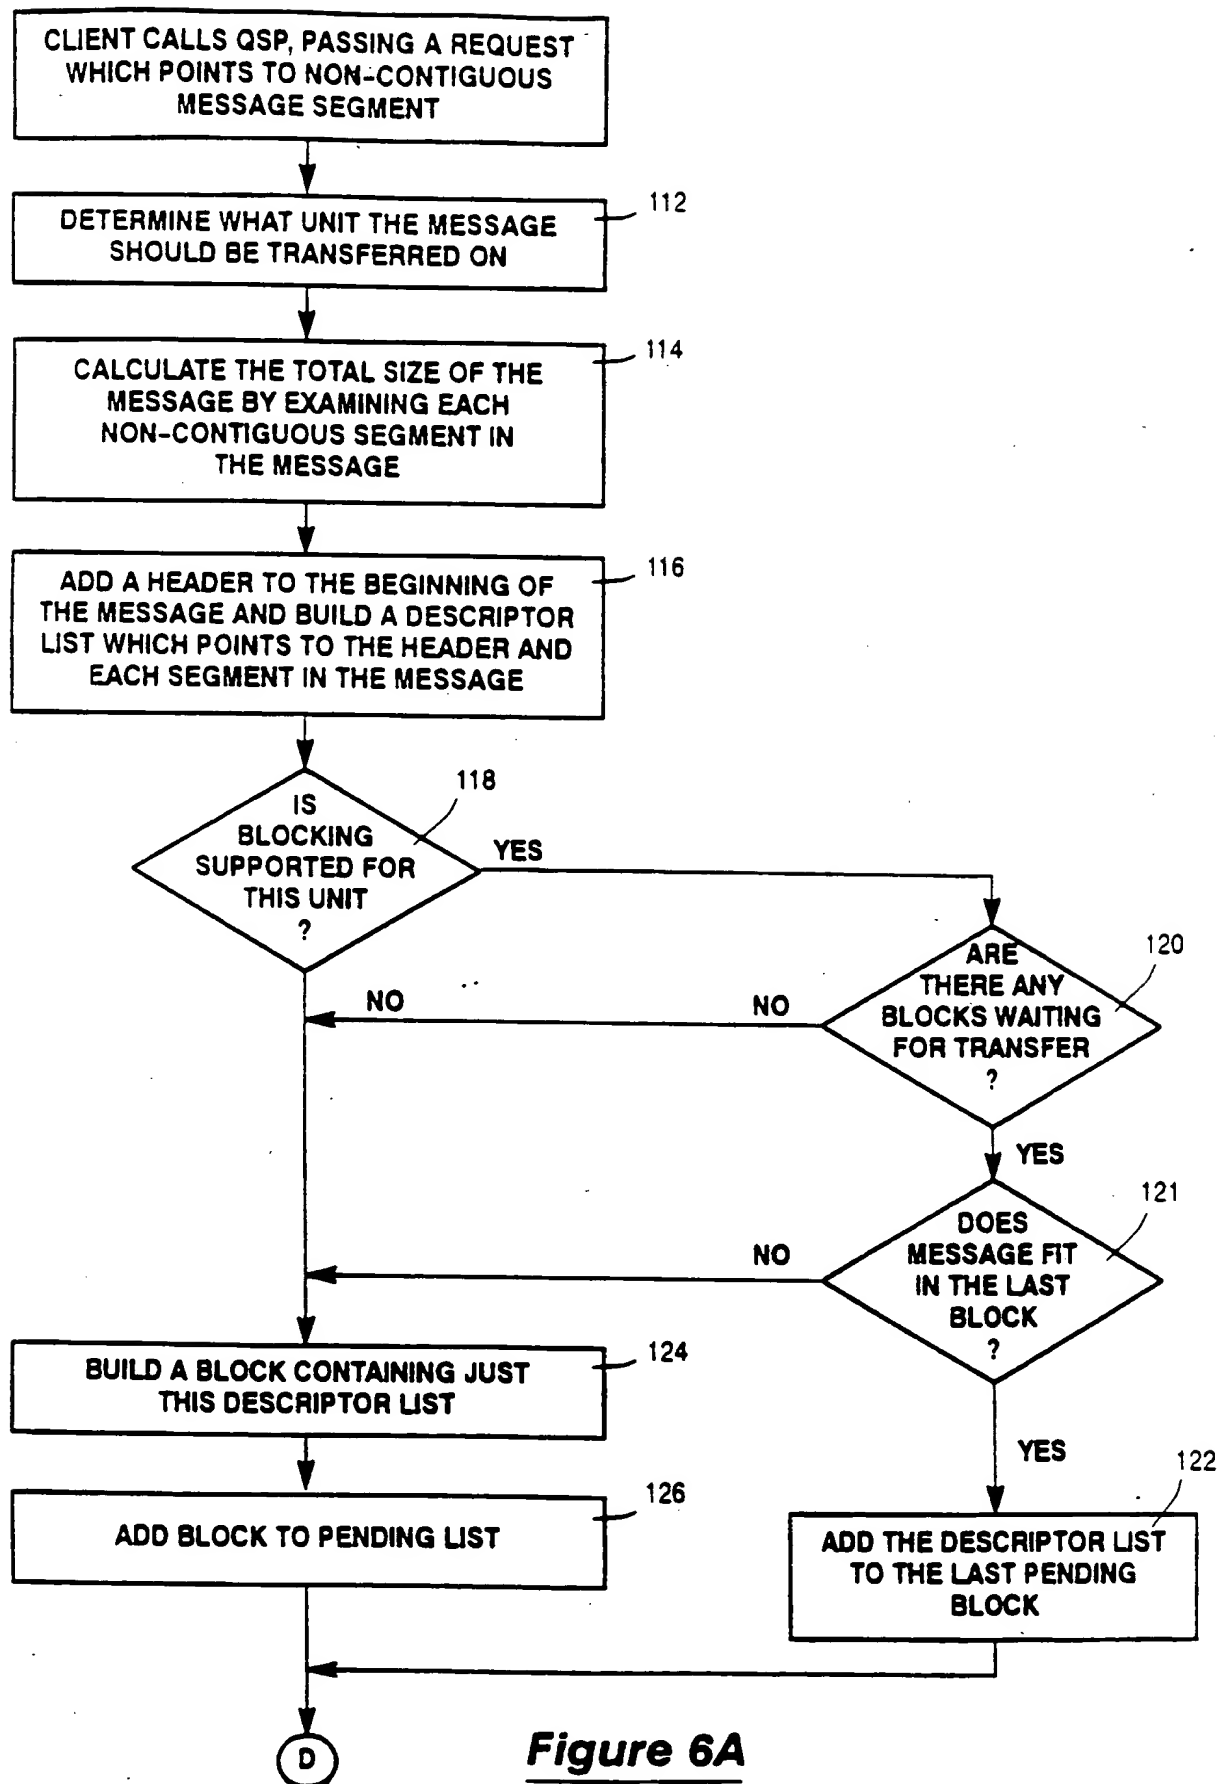



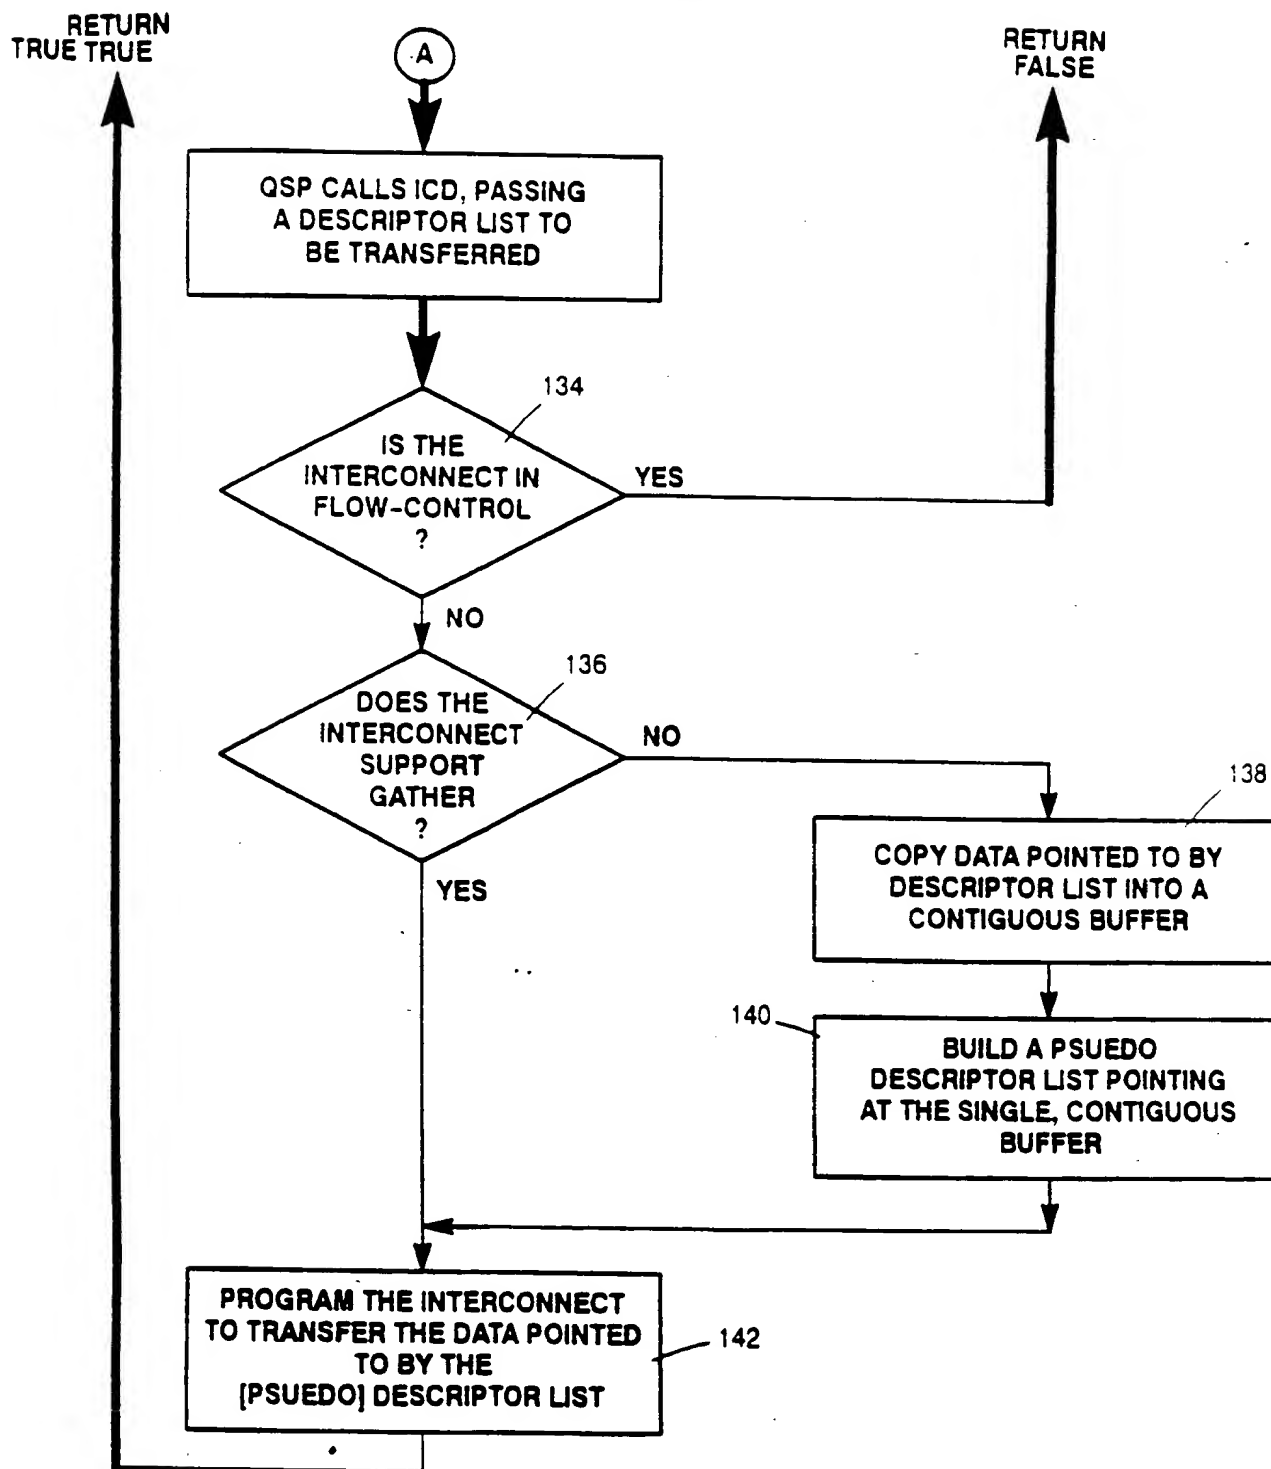
Figure 5

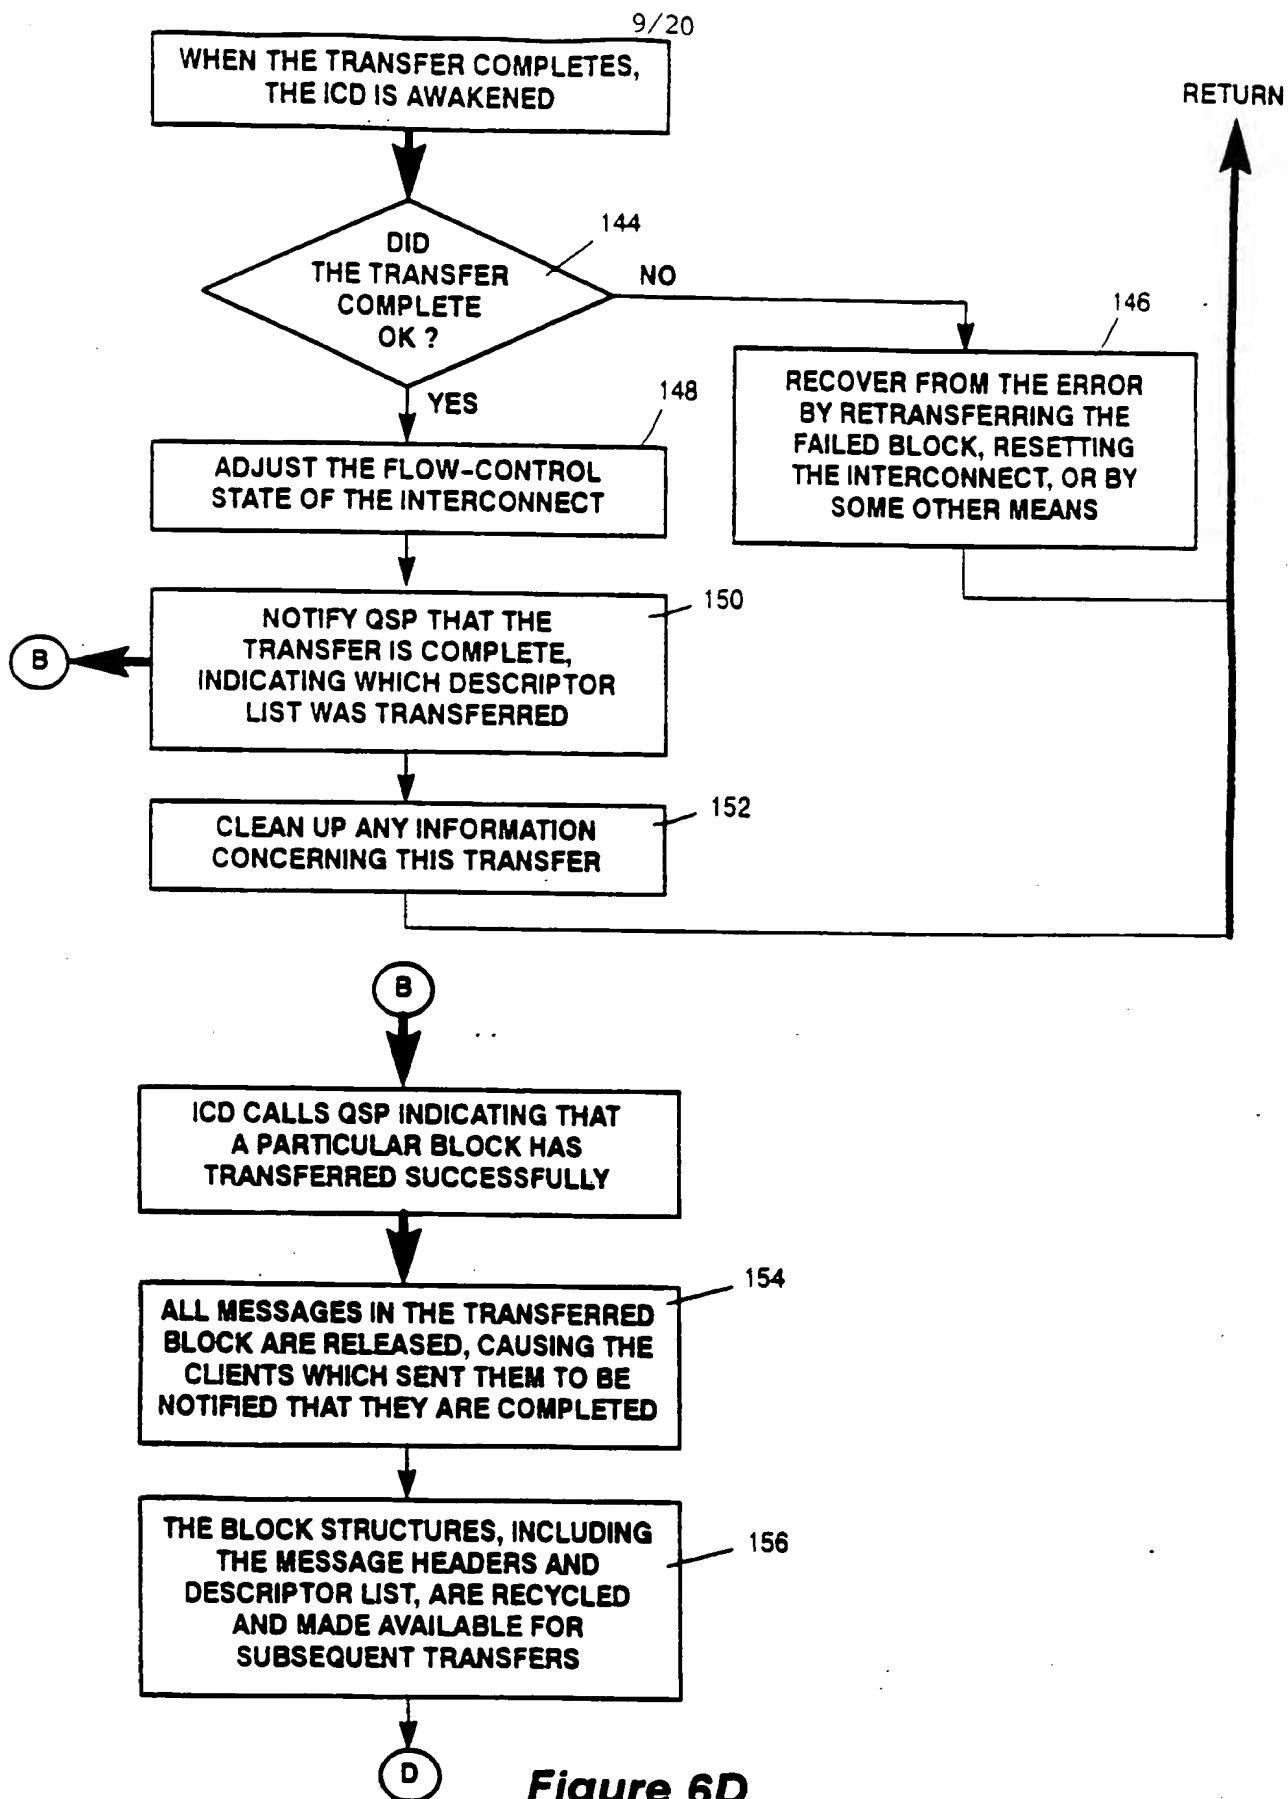
**Figure 6A**

7/20

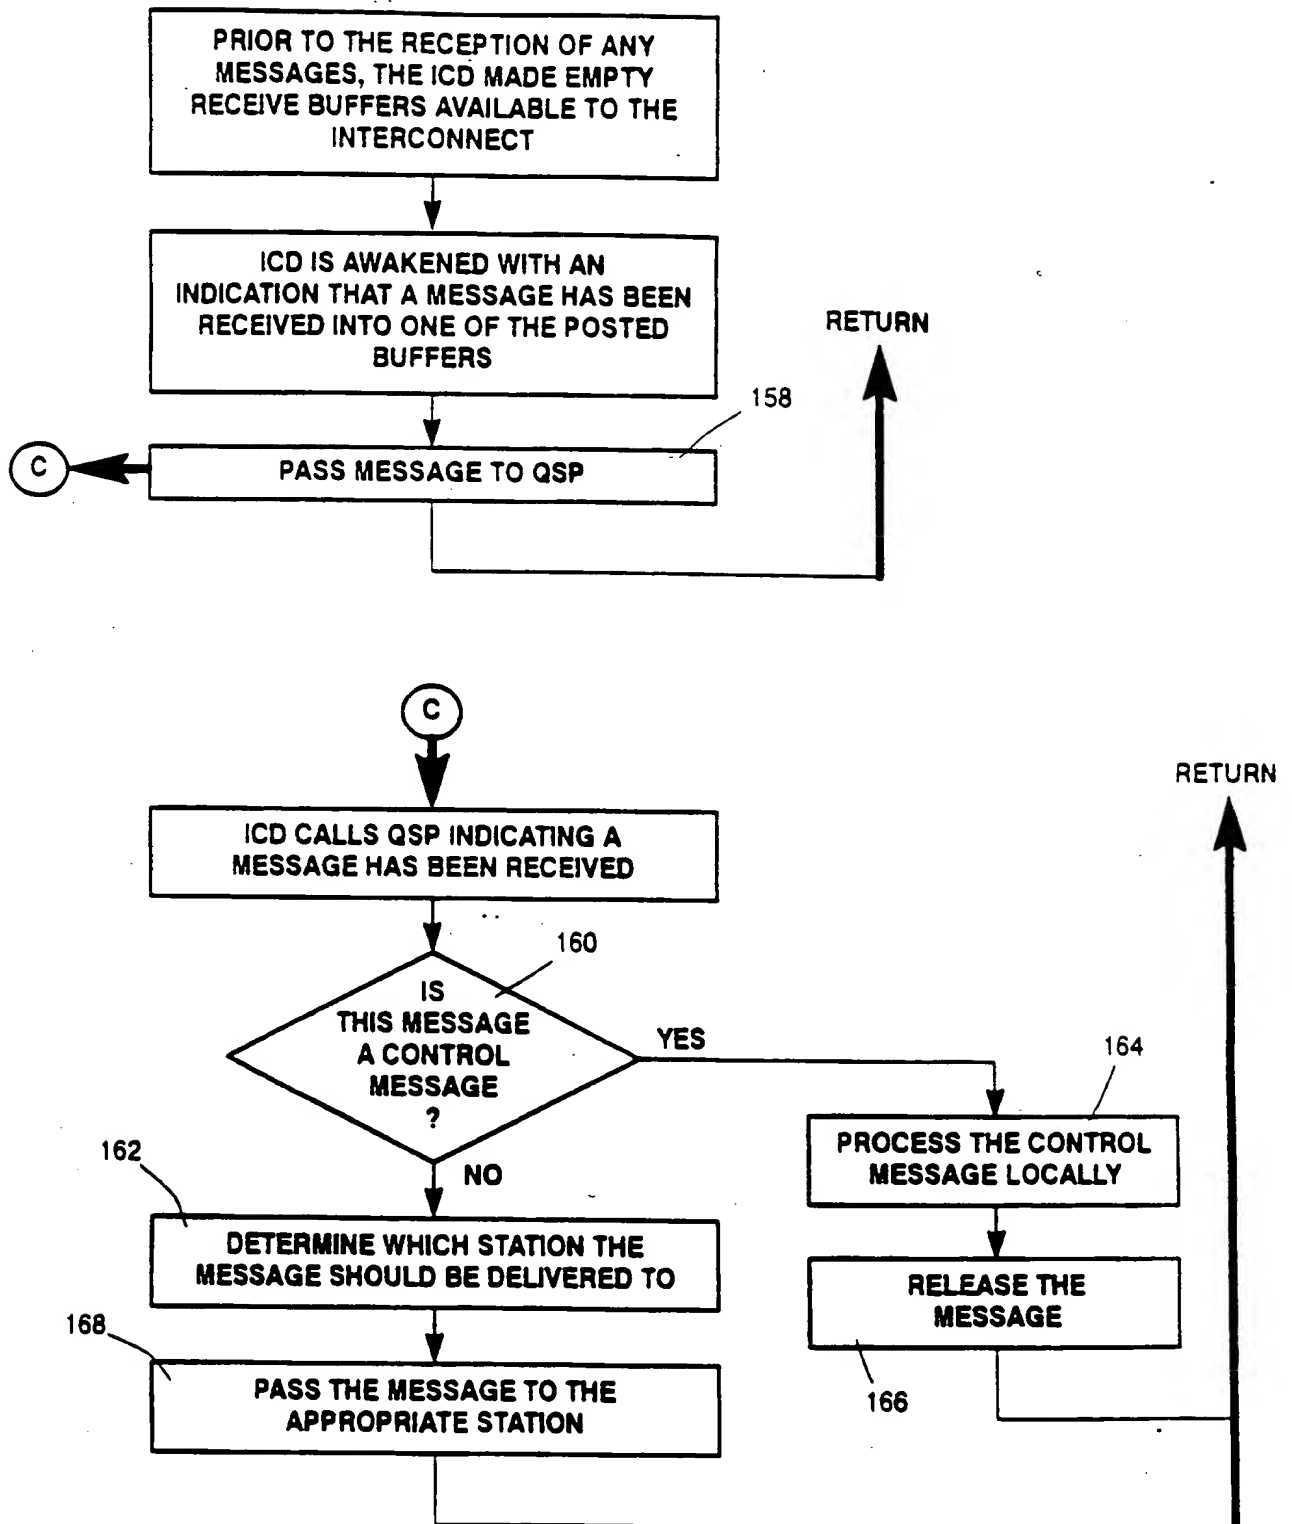
**Figure 6B**

3/20

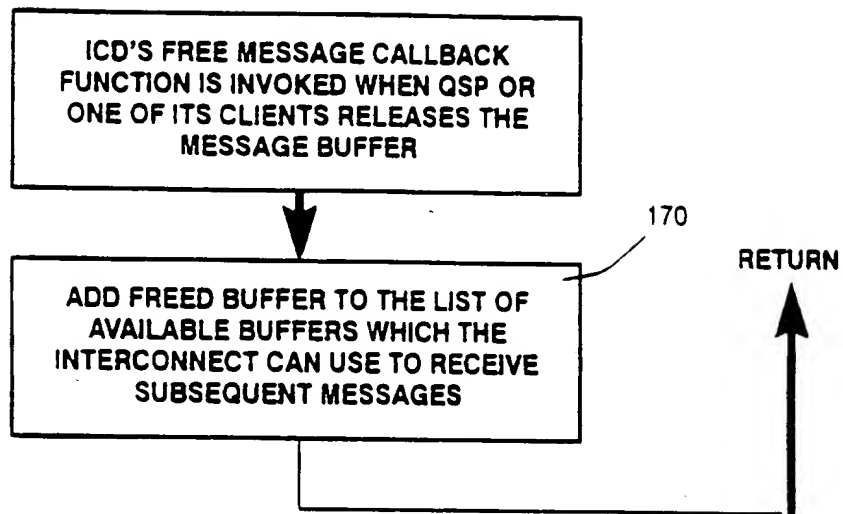
**Figure 6C**

**Figure 6D**


10/20

**Figure 6E**

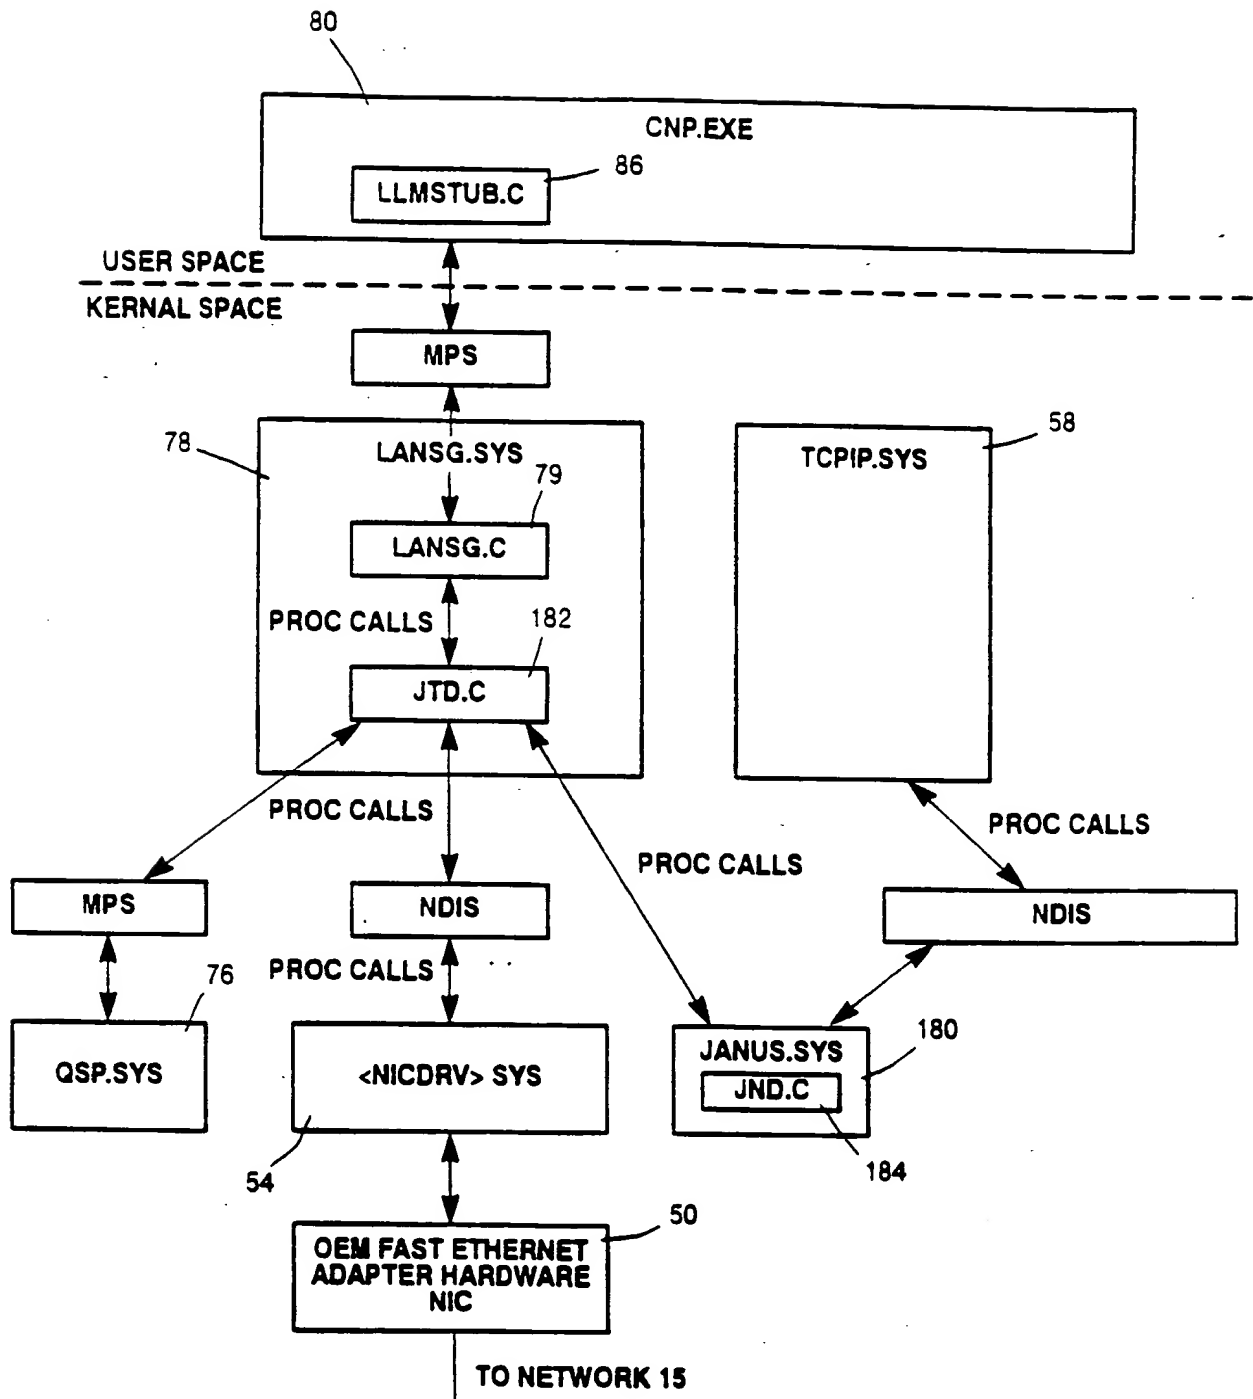
11/20

**Figure 6F**

12/20

**Figure 7**

13/20

**Figure 8**

14/20

```

if (rcvd_frame_ptr->MAC_hdr.length <= 1500) { 200
    /* frame is an 802.3 frame */
    if (rcvd_frame_ptr->LLC_hdr.DSAP == 0xAA) { 202
        /* frame is a SNAP frame */
        if (rcvd_frame_ptr->LLC_hdr.control == 0x03) { 204
            /* frame is a UI frame */
            switch (rcvd_frame_ptr->SNAP_hdr.type) {
                case 0x0800:
                    type_of_frame_rcvd = IP;
                    break;
                case 0x0806:
                    type_of_frame_rcvd = ARP;
                    break;
                default:
                    type_of_frame_rcvd = UNKNOWN;
                    break;
            }
        } else {
            type_of_frame_rcvd = UNKNOWN; 208
        }
    } else if (rcvd_frame_ptr->LLC_hdr.DSAP == 0x04) {
        if (lookup_source_addr(rcvd_frame_ptr->MAC_hdr.source_addr, &src_ptr)) { 210
            type_of_frame_rcvd = BNA; 212
        } else {
            type_of_frame_rcvd = UNKNOWN;
        }
    } else {
        type_of_frame_rcvd = UNKNOWN;
    }
} else {
    /* frame is either an Ethernet frame or a large BNA frame */
    /* find out if the frame is a large BNA frame */

```

Figure 9A

15/20

```

if (rcvd_frame_ptr->LLC_hdr.DSAP == 0x04) { 216
    if (lookup_source_addr(rcvd_frame_ptr->MAC_hdr.source_addr, &stn_ptr)) { 218
        type_of_frame_rcvd = BNA; 220
    } else {
        type_of_frame_rcvd = UNKNOWN; 222
    }
} else {
    /* must be an IP or ARP frame */
    switch (rcvd_frame_ptr->MAC_hdr.type) {
        case 0x0800:
            type_of_frame_rcvd = IP;
            break;
        case 0x0806:
            type_of_frame_rcvd = ARP;
            break;
        default:
            type_of_frame_rcvd = UNKNOWN;
            break;
    }
}

switch (type_of_frame_rcvd) {
    case IP:
    case ARP:
        /* determine routing of frame based on destination IP address */ 226
        if (rcvd_frame_ptr->IP_hdr.dest_addr == IP_BROADCAST_ADDR) {
            ip_addr_owner = A Series & NT;
        } else {
            /* first, assume that IP address is owned by neither A Series or NT */
            ip_addr_owner = NEITHER; 228
            /* determine if IP address is either the A Series Server's and/or the

```

Figure 9B

16/20

```

    NT Server's multicast address list */
    if (rcvd_frame_ptr->IP_hdr.dest_addr "in" A Series multicast_addr_list)
230      ip_addr_owner = A Series;
    if (rcvd_frame_ptr->IP_hdr.dest_addr "in" NT multicast_addr_list)
      ip_addr_owner &= NT;
  }
  /* route the frame */
  if (ip_addr_owner == A Series) {
    /* send it up to LANSg for forwarding to TCP/IP on A Series server */
  } else if (ip_addr_owner == NT) {
    /* forwarding it to NT's TCPIP.SYS */
232  } else {
    /* determine who to send it to based on the values of
       EnableIPRoutingInMCP and EnableIPRouting in Registry */
  }
  break;

234  case BNA:
    /* send it directly to the LANSg passing the correct stn_ptr */
    break;

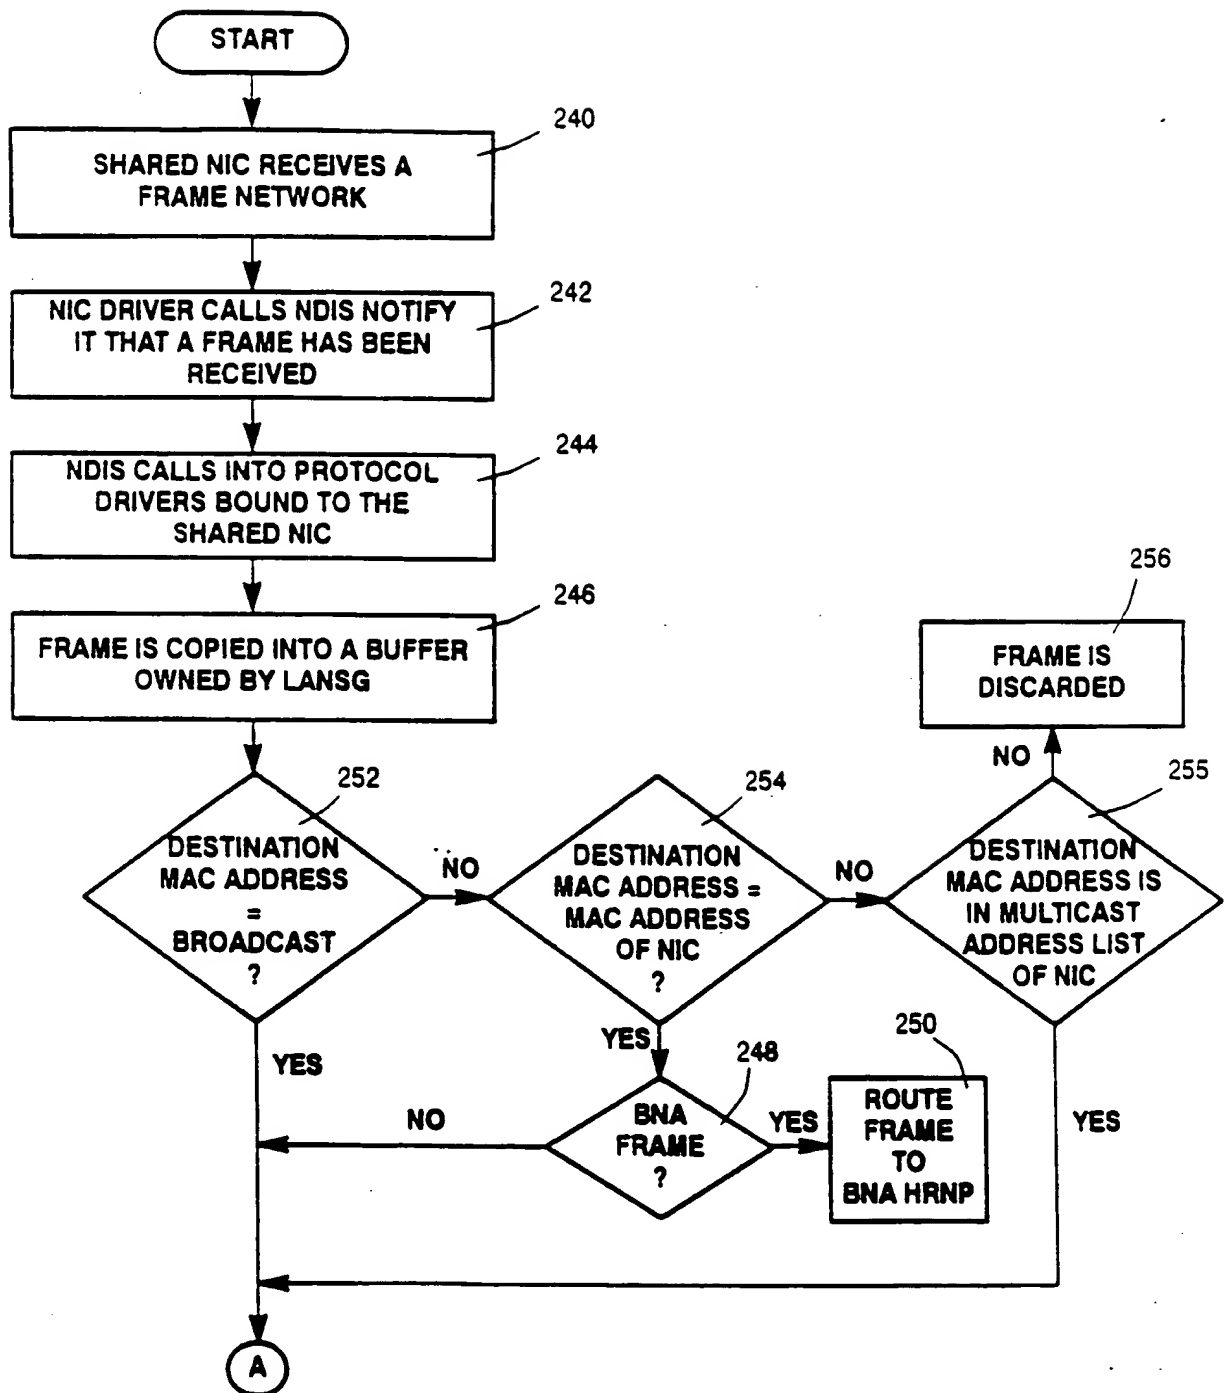
236  case UNKNOWN:
    /* forward it up through NDIS */
    break;

  default:
    /* shouldn't happen */
    break;
}

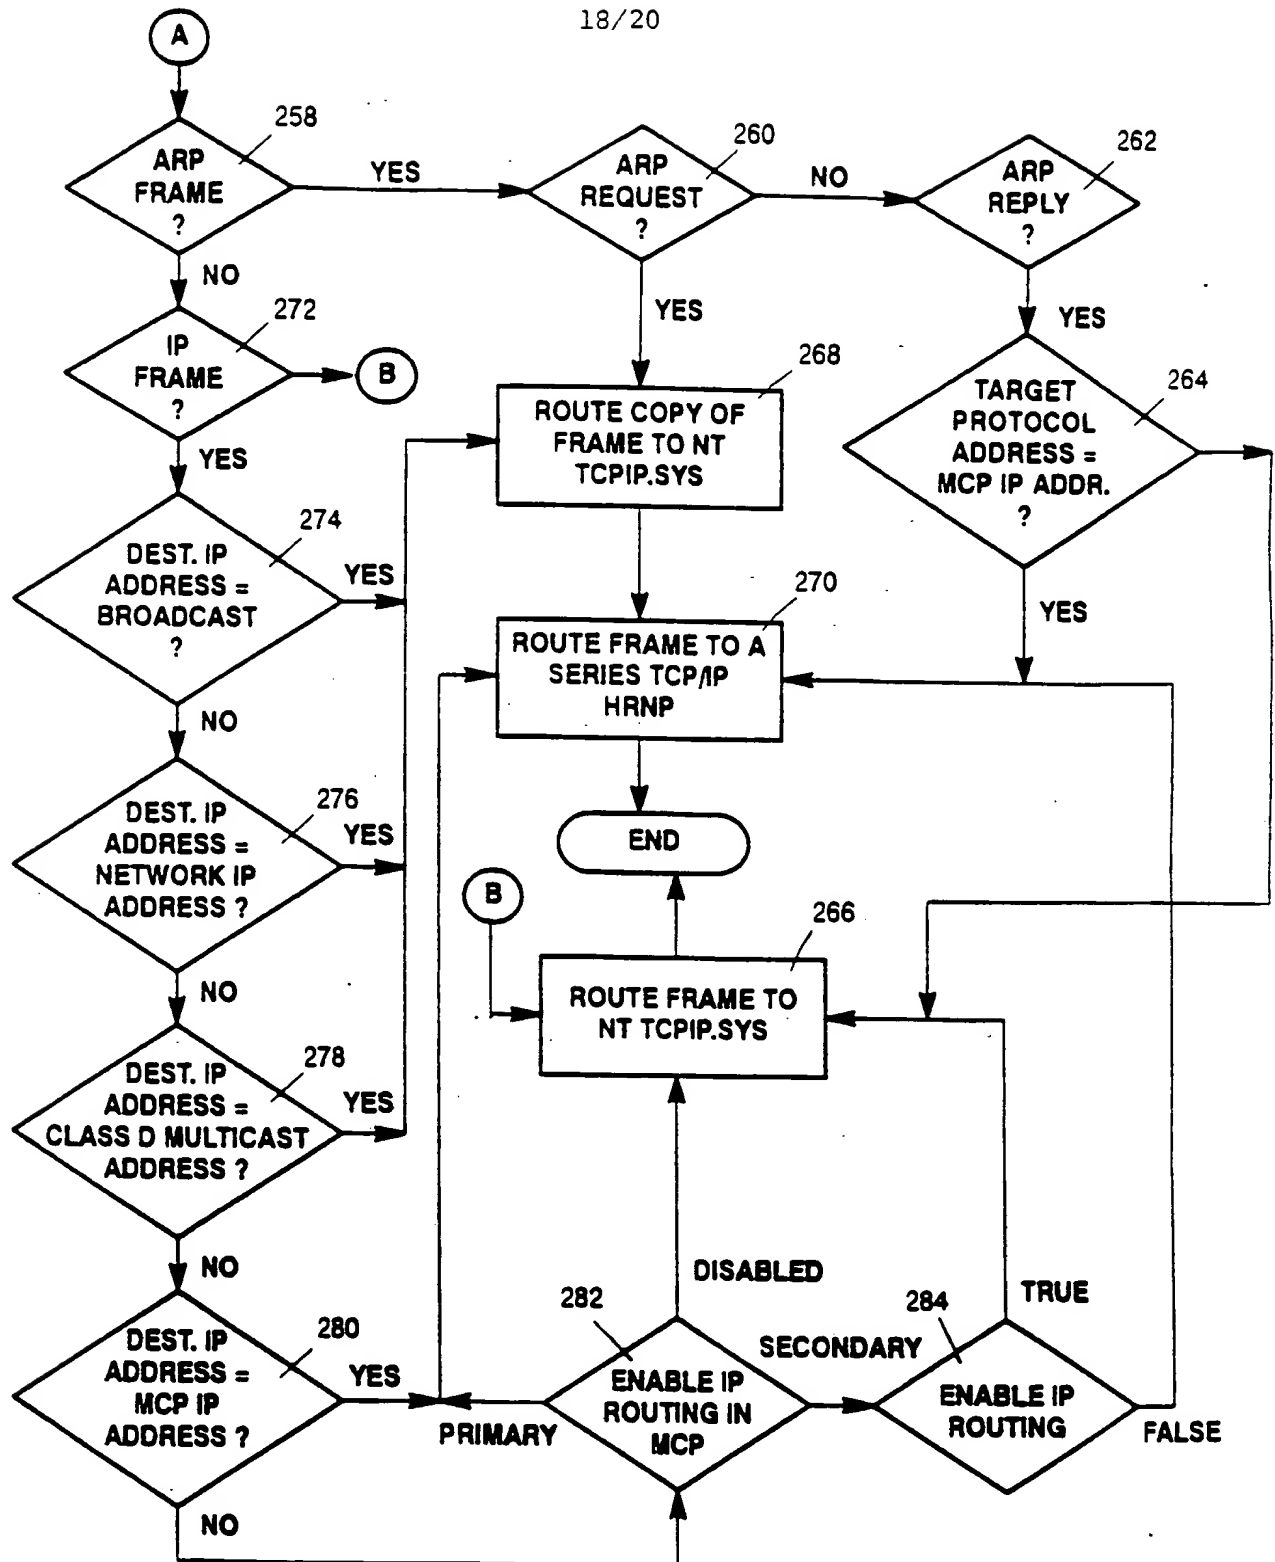
```

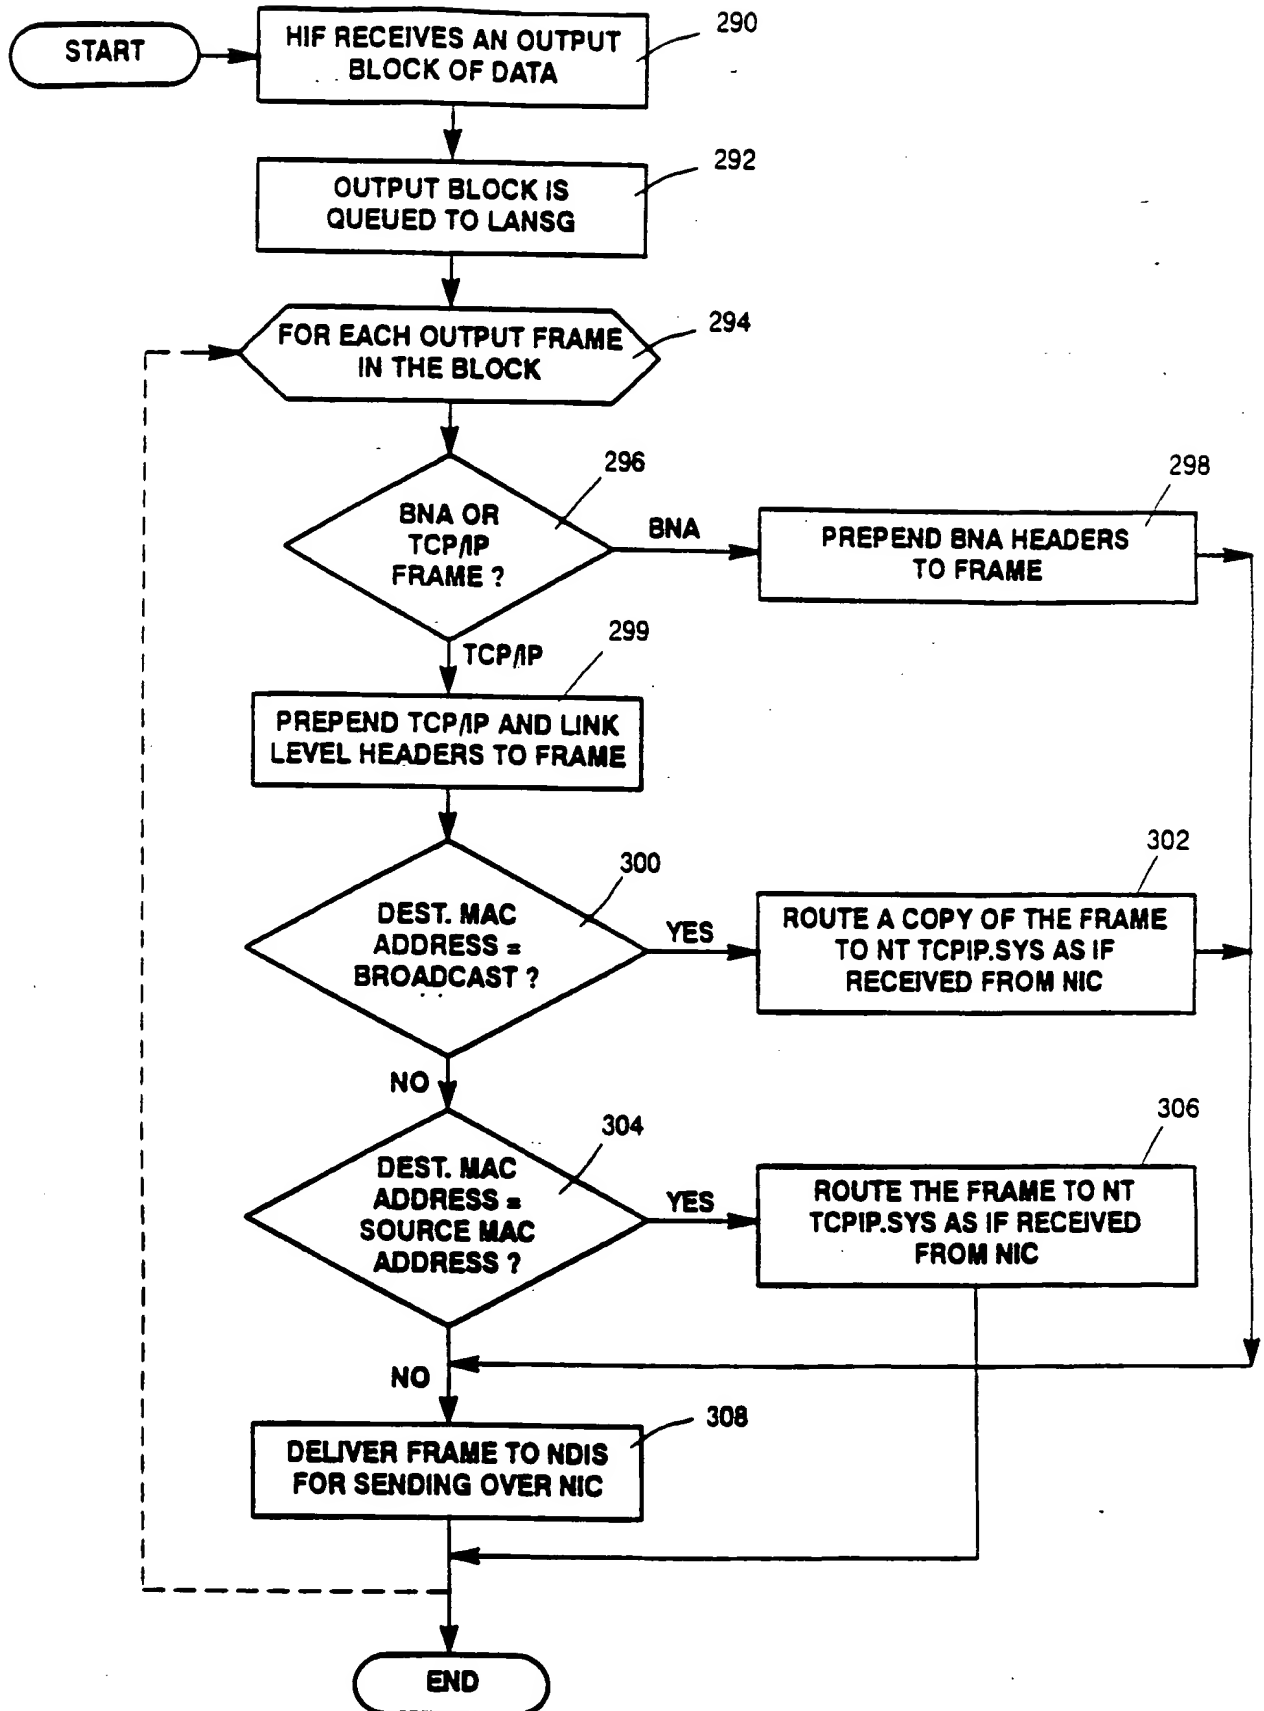
Figure 9C

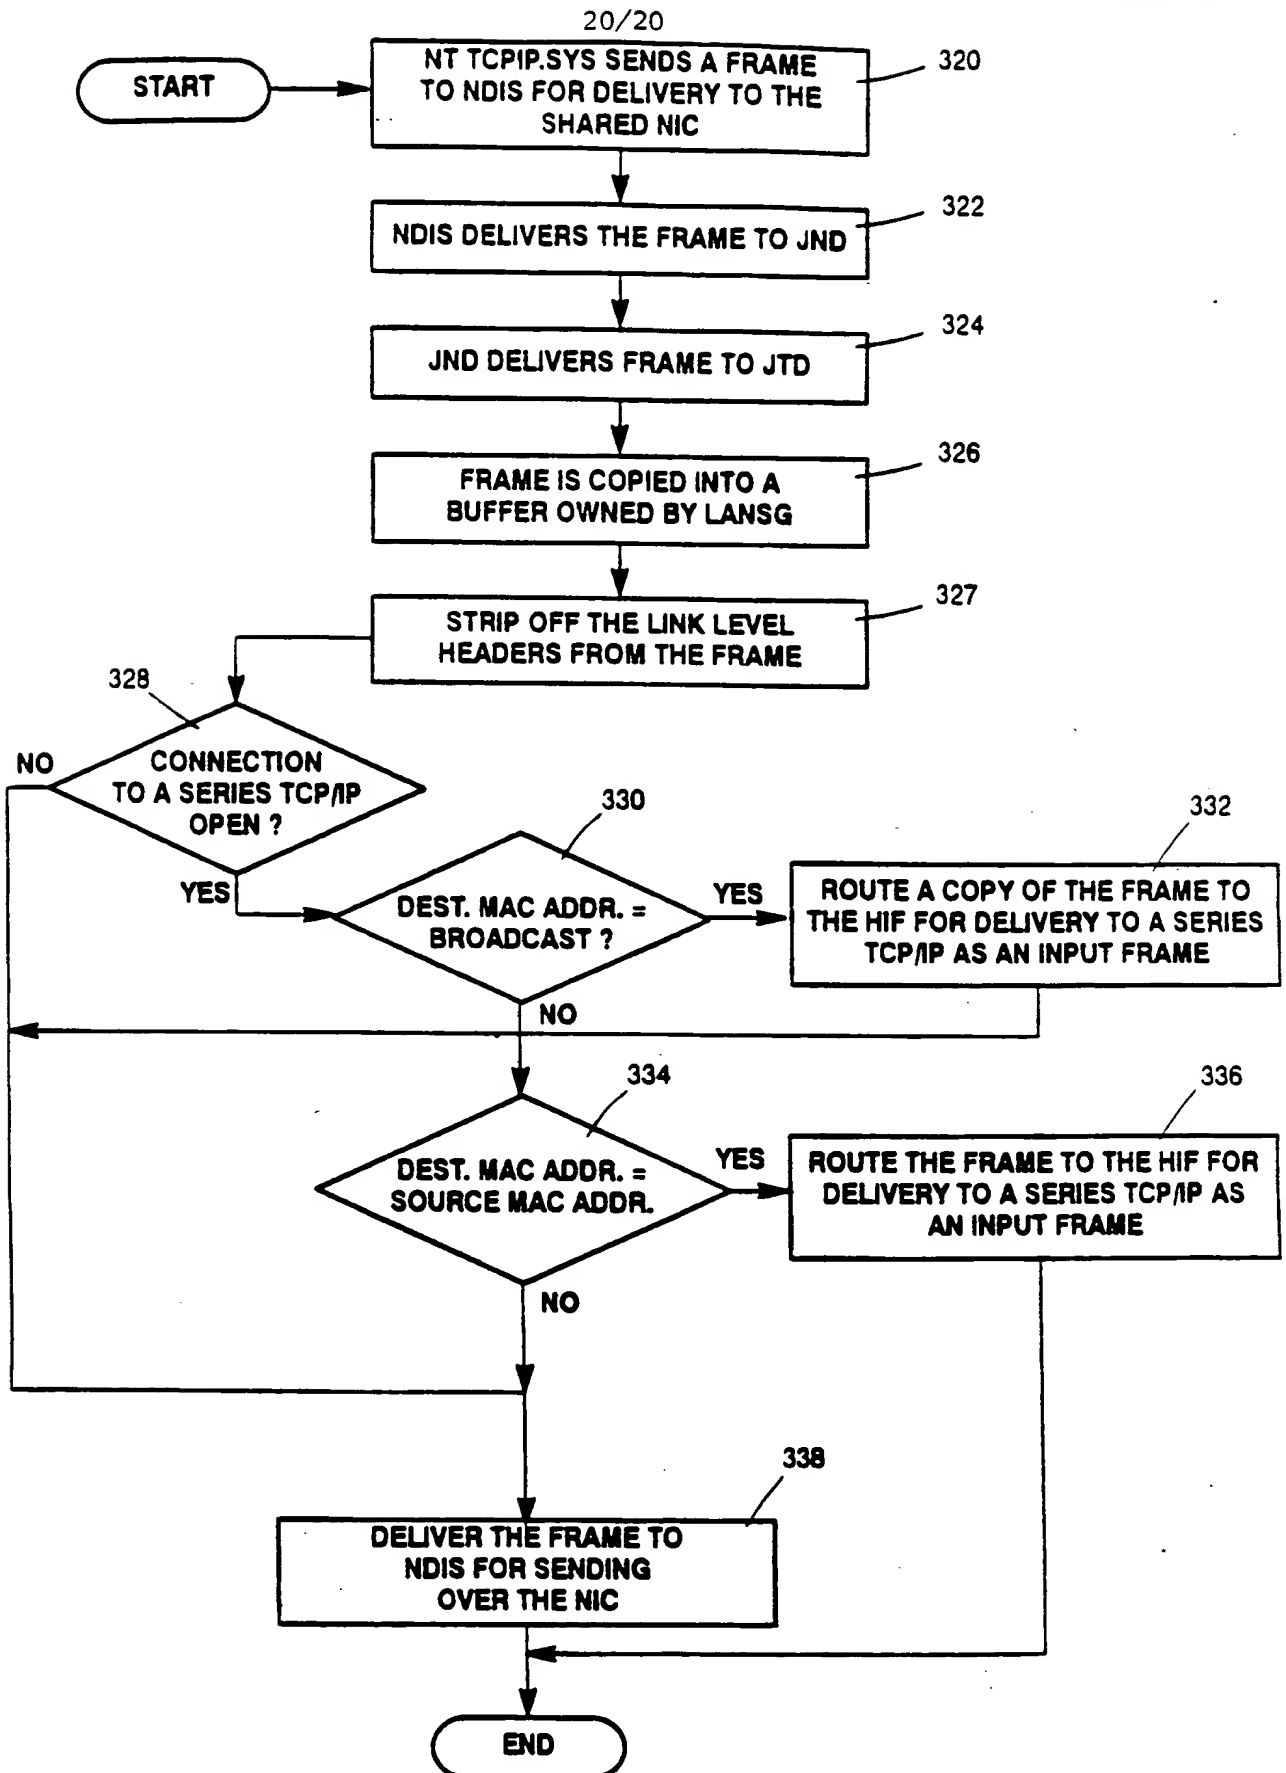
17/20

**Figure 10A**

18/20

**Figure 10B**

**Figure 10C**

**Figure 10D**

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/US 98/11201

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04L29/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>WO 97 01944 A (BULL HN INFORMATION SYST) 16 January 1997 see page 4, line 24 - page 8, line 2 see page 9, line 29 - page 10, line 8 see page 13, line 6 - page 18, line 20 see page 26, line 9 - page 27, line 29 -----</p>	1, 20

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but after the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"Z" document member of the same patent family

Date of the actual completion of the international search

6 November 1998

Date of mailing of the international search report

13/11/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Dupuis, H